AMD Accelerated Parallel Processing

OpenCL™ User Guide

October 2014

rev1.0

© 2013 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, AMD Accelerated Parallel Processing, the AMD Accelerated Parallel Processing logo, ATI, the ATI logo, Radeon, FireStream, FirePro, Catalyst, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Microsoft, Visual Studio, Windows, and Windows Vista are registered trademarks of Microsoft Corporation in the U.S. and/or other jurisdictions. Other names are for informational purposes only and may be trademarks of their respective owners. OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.

The contents of this document are provided in connection with Advanced Micro Devices, Inc. ("AMD") products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. The information contained herein may be of a preliminary or advance nature and is subject to change without notice. No license, whether express, implied, arising by estoppel or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD's Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD's products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD's product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.

Advanced Micro Devices, Inc.
One AMD Place
P.O. Box 3453
Sunnyvale, CA 94088-3453
www.amd.com

For AMD Accelerated Parallel Processing:

URL:	**developer.amd.com/appsdk**
Developing:	**developer.amd.com/**
Forum:	**developer.amd.com/openclforum**

Preface

About This Document

This document provides a basic description of the AMD Accelerated Parallel Processing environment and components. It describes the basic architecture of compute and stream processors. This document also provides a guide for programmers who want to use AMD Accelerated Parallel Processing to accelerate their applications.

Audience

This document is intended for programmers. It assumes prior experience in writing code for CPUs and a basic understanding of threads (work-items). While a basic understanding of GPU architectures is useful, this document does not assume prior graphics knowledge. It further assumes an understanding of chapters 1, 2, and 3 of the *OpenCL Specification* (for the latest version, see http://www.khronos.org/registry/cl/).

Organization

This AMD Accelerated Parallel Processing document begins, in Chapter 1, with an overview of: the AMD Accelerated Parallel Processing programming models, OpenCL, and the AMD Compute Abstraction Layer (CAL). Chapter 2 discusses the AMD implementation of OpenCL. Chapter 3 discusses the compiling and running of OpenCL programs. Chapter 4 describes using the AMD CodeXL GPU Debugger and the GNU debugger (GDB) to debug OpenCL programs. Chapter 5 provides information about the extension that defines the OpenCL Static C++ kernel language, which is a form of the ISO/IEC Programming languages C++ specification. Appendix A describes the supported optional OpenCL extensions. Appendix B details the installable client driver (ICD) for OpenCL. Appendix C details the compute kernel and contrasts it with a pixel shader. Appendix C describes the OpenCL binary image format (BIF). Appendix D provides a hardware overview of pre-GCN devices. Appendix E describes the interoperability between OpenCL and OpenGL. Appendix F describes the new and deprecated functions in OpenCL 2.0. The last section of this book is an index.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Conventions

The following conventions are used in this document.

`mono-spaced font`	A filename, file path, or code.
*	Any number of alphanumeric characters in the name of a code format, parameter, or instruction.
[1,2)	A range that includes the left-most value (in this case, 1) but excludes the right-most value (in this case, 2).
[1,2]	A range that includes both the left-most and right-most values (in this case, 1 and 2).
{x | y}	One of the multiple options listed. In this case, x or y.
0.0f 0.0	A single-precision (32-bit) floating-point value. A double-precision (64-bit) floating-point value.
1011b	A binary value, in this example a 4-bit value.
7:4	A bit range, from bit 7 to 4, inclusive. The high-order bit is shown first.
italicized word or phrase	The first use of a term or concept basic to the understanding of stream computing.

Related Documents

- *The OpenCL Specification*, Version 1.1, Published by Khronos OpenCL Working Group, Aaftab Munshi (ed.), 2010.
- *The OpenCL Specification*, Version 2.0, Published by Khronos OpenCL Working Group, Aaftab Munshi (ed.), 2013.
- AMD, *R600 Technology, R600 Instruction Set Architecture*, Sunnyvale, CA, est. pub. date 2007. This document includes the RV670 GPU instruction details.
- ISO/IEC 9899:TC2 - *International Standard - Programming Languages - C*
- Kernighan Brian W., and Ritchie, Dennis M., *The C Programming Language*, Prentice-Hall, Inc., Upper Saddle River, NJ, 1978.
- I. Buck, T. Foley, D. Horn, J. Sugerman, K. Fatahalian, M. Houston, and P. Hanrahan, *"Brook for GPUs: stream computing on graphics hardware,"* ACM Trans. Graph., vol. 23, no. 3, pp. 777–786, 2004.
- *AMD Compute Abstraction Layer (CAL) Intermediate Language (IL) Reference Manual*. Published by AMD.
- Buck, Ian; Foley, Tim; Horn, Daniel; Sugerman, Jeremy; Hanrahan, Pat; Houston, Mike; Fatahalian, Kayvon. "BrookGPU" http://graphics.stanford.edu/projects/brookgpu/
- Buck, Ian. "Brook Spec v0.2". October 31, 2003. http://merrimac.stanford.edu/brook/brookspec-05-20-03.pdf
- *OpenGL Programming Guide*, at http://www.glprogramming.com/red/
- *Microsoft DirectX Reference Website*, at http://msdn.microsoft.com/en-us/directx

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

- *GPGPU*: http://www.gpgpu.org, and Stanford BrookGPU discussion forum http://www.gpgpu.org/forums/

Contact Information

URL:	**developer.amd.com/appsdk**
Developing:	**developer.amd.com/**
Forum:	**developer.amd.com/openclforum**

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Contents

Preface

Contents

Chapter 1

OpenCL Architecture and AMD Accelerated Parallel Processing

1.1 Terminology 1-1
1.2 OpenCL Overview 1-2
1.3 Programming Model 1-3
1.4 Synchronization 1-4
1.5 Memory Architecture and Access 1-5
1.5.1 Data Share Operations 1-7
1.5.2 Dataflow in Memory Hierarchy 1-8
1.5.3 Memory Access 1-9
1.5.4 Global Memory 1-10
1.5.5 Image Read/Write 1-10
1.6 Example Programs 1-10
1.6.1 First Example: Simple Buffer Write 1-10
1.6.2 Example: SAXPY Function 1-14
1.6.3 Example: Parallel Min() Function 1-19

Chapter 2

AMD Implementation

2.1 The AMD Accelerated Parallel Processing Implementation of OpenCL 2-1
2.1.1 Work-Item Processing 2-4
2.1.2 Work-Item Creation 2-4
2.1.3 Flow Control 2-4
2.2 Hardware Overview for GCN Devices 2-6
2.2.1 Key differences between pre-GCN and GCN devices 2-7
2.2.2 Key differences between Southern Islands, Sea Islands, and Volcanic Islands families 2-8
2.3 Communication Between Host and the GPU Compute Device 2-8
2.3.1 Processing API Calls: The Command Processor 2-9
2.3.2 DMA Transfers 2-9
2.3.3 Masking Visible Devices 2-10
2.4 Wavefront Scheduling 2-10

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Chapter 3
Building and Running OpenCL Programs
3.1 Compiling the Program ... 3-2
3.1.1 Compiling on Windows ... 3-2
3.1.2 Compiling on Linux ... 3-3
3.1.3 Supported Standard OpenCL Compiler Options ... 3-4
3.1.4 AMD-Developed Supplemental Compiler Options ... 3-4
3.2 Running the Program ... 3-5
3.2.1 Running Code on Windows ... 3-6
3.2.2 Running Code on Linux ... 3-7
3.3 Calling Conventions ... 3-7

Chapter 4
Debugging OpenCL
4.1 AMD CodeXL GPU Debugger ... 4-1
4.2 Debugging CPU Kernels with GDB ... 4-2
4.2.1 Setting the Environment ... 4-2
4.2.2 Setting the Breakpoint in an OpenCL Kernel ... 4-2
4.2.3 Sample GDB Session ... 4-3
4.2.4 Notes ... 4-4

Chapter 5
OpenCL Static C++ Programming Language
5.1 Overview ... 5-1
5.1.1 Supported Features ... 5-1
5.1.2 Unsupported Features ... 5-2
5.1.3 Relations with ISO/IEC C++ ... 5-2
5.2 Additions and Changes to Section 5 - The OpenCL C Runtime ... 5-2
5.2.1 Additions and Changes to Section 5.7.1 - Creating Kernel Objects ... 5-2
5.2.2 Passing Classes between Host and Device ... 5-3
5.3 Additions and Changes to Section 6 - The OpenCL C Programming Language ... 5-3
5.3.1 Building C++ Kernels ... 5-3
5.3.2 Classes and Derived Classes ... 5-3
5.3.3 Namespaces ... 5-4
5.3.4 Overloading ... 5-4
5.3.5 Templates ... 5-5
5.3.6 Exceptions ... 5-6
5.3.7 Libraries ... 5-6
5.3.8 Dynamic Operation ... 5-6
5.4 Examples ... 5-6
5.4.1 Passing a Class from the Host to the Device and Back ... 5-6
5.4.2 Kernel Overloading ... 5-7
5.4.3 Kernel Template ... 5-8

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Chapter 6

OpenCL 2.0

6.1 Introduction 6-1
6.2 Shared Virtual Memory (SVM) 6-1
6.2.1 Overview 6-1
6.2.2 Usage 6-3
Coarse-grained memory 6-4
6.3 Generic Address Space 6-5
6.3.1 Overview 6-5
6.3.2 Usage 6-6
Generic example 6-6
AMD APP SDK example 6-6
6.4 Device-side enqueue 6-7
6.4.1 Overview 6-7
6.4.2 Usage 6-8
Iterate until convergence 6-8
Data-dependent refinement 6-8
Extracting Primes from an array by using device-side enqueue 6-9
Binary search using device enqueue (or kernel enqueue) 6-15
6.5 Atomics and synchronization 6-18
6.5.1 Overview 6-18
6.5.2 Usage 6-18
Atomic Loads/Stores 6-19
Atomic Compare and Exchange (CAS) 6-20
Atomic Fetch 6-20
6.6 Pipes 6-21
6.6.1 Overview 6-21
6.6.2 Functions for accessing pipes 6-22
6.6.3 Usage 6-23
6.7 Sub-groups 6-24
6.7.1 Overview 6-24
6.8 Program-scope global Variables 6-24
6.8.1 Overview 6-24
6.9 Image Enhancements 6-24
6.9.1 Overview 6-24
6.9.2 sRGB 6-25
6.9.3 Depth images 6-27
6.10 Non-uniform work group size 6-28
6.10.1 Overview 6-28
6.11 Portability considerations 6-28
6.11.1 Migrating from OpenCL 1.2 to OpenCL 2.0 6-28
6.11.2 Identifying implementation specifics 6-29

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Appendix A
OpenCL Optional Extensions

A.1 Extension Name Convention A-1
A.2 Querying Extensions for a Platform A-1
A.3 Querying Extensions for a Device A-2
A.4 Using Extensions in Kernel Programs A-2
A.5 Getting Extension Function Pointers A-3
A.6 List of Supported Extensions that are Khronos-Approved A-3
A.7 cl_ext Extensions A-4
A.8 AMD Vendor-Specific Extensions A-4
A.8.1 cl_amd_fp64 ..... A-4
A.8.2 cl_amd_vec3 ..... A-4
A.8.3 cl_amd_device_persistent_memory A-4
A.8.4 cl_amd_device_attribute_query A-5
cl_device_profiling_timer_offset_amd A-5
cl_amd_device_topology A-5
cl_amd_device_board_name A-5
A.8.5 cl_amd_compile_options A-6
A.8.6 cl_amd_offline_devices A-6
A.8.7 cl_amd_event_callback A-6
A.8.8 cl_amd_popcnt A-7
A.8.9 cl_amd_media_ops A-7
A.8.10 cl_amd_media_ops2 ..... A-9
A.8.11 cl_amd_printf A-12
A.8.12 cl_amd_predefined_macros A-13
A.8.13 cl_amd_bus_addressable_memory A-14
A.9 Supported Functions for cl_amd_fp64 / cl_khr_fp64 ..... A-15
A.10 Extension Support by Device A-15

Appendix B
The OpenCL Installable Client Driver (ICD)

B.1 Overview B-1
B.2 Using ICD B-1

Appendix C
OpenCL Binary Image Format (BIF) v2.0

C.1 Overview C-1
C.1.1 Executable and Linkable Format (ELF) Header C-2
C.1.2 Bitness C-3
C.2 BIF Options C-3

Appendix D

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Hardware overview of pre-GCN devices

Appendix E
OpenCL-OpenGL Interoperability

E.1 Under Windows E-1
E.1.1 Single GPU Environment E-2
Creating CL Context from a GL Context E-2
E.1.2 Multi-GPU Environment E-4
Creating CL context from a GL context E-4
E.1.3 Limitations E-7
E.2 Linux Operating System E-7
E.2.1 Single GPU Environment E-7
Creating CL Context from a GL Context E-7
E.2.2 Multi-GPU Configuration E-10
Creating CL Context from a GL Context E-10
E.3 Additional GL Formats Supported E-13

Appendix F
New and deprecated functions in OpenCL 2.0

F.1 New built-in functions in OpenCL 2.0 ..... F-1
F.1.1 Work Item Functions F-1
F.1.2 Integer functions F-1
F.1.3 Synchronization Functions F-1
F.1.4 Address space quailfier functions F-2
F.1.5 Atomic functions F-2
F.1.6 Image Read and Write Functions F-2
F.1.7 Work group functions F-2
F.1.8 Pipe functions F-3
F.1.9 Enqueueing Kernels F-3
F.1.10 Sub-groups F-4
F.2 Deprecated built-ins F-4
F.3 New runtime APIs in OpenCL 2.0 ..... F-6
F.3.1 New Types F-6
F.3.2 New Macros F-6
F.3.3 New API calls F-8
F.4 Deprecated runtimes F-8

Index

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Figures

1.1 OpenCL Programming Model 1-3
1.2 Interrelationship of Memory Domains for Southern Islands Devices 1-6
1.3 Dataflow between Host and GPU 1-6
1.4 High-Level Memory Configuration 1-7
1.5 Memory Hierarchy Dataflow 1-8
2.1 AMD Accelerated Parallel Processing Software Ecosystem 2-1
2.2 Simplified Mapping of OpenCL onto AMD Accelerated Parallel Processing 2-2
2.3 Work-Item Grouping Into Work-Groups and Wavefronts 2-3
2.4 Generalized AMD GPU Compute Device Structure for GCN Devices 2-6
2.5 AMD Radeon™ HD 79XX Device Partial Block Diagram 2-7
2.6 Simplified Execution Of Wavefront On A Single Compute Unit 2-11
2.7 Compute Unit Stall Due to Data Dependency 2-12
3.1 OpenCL Compiler Toolchain 3-1
3.2 Runtime Processing Structure 3-6
A.1 Peer-to-Peer Transfers Using the `cl_amd_bus_addressable_memory` Extension A-14
D.1 Generalized AMD GPU Compute Device Structure D-2
D.2 Simplified Block Diagram of an Evergreen-Family GPU D-3

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Tables

A.1 Extension Support for AMD GPU Devices 1 ..... A-15
A.2 Extension Support for Older AMD GPUs and CPUs A-16
C.1 ELF Header Fields C-2
E.1 AMD-Supported GL Formats E-14

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Chapter 1
OpenCL Architecture and AMD Accelerated Parallel Processing

This chapter provides a general software and hardware overview of the AMD Accelerated Parallel Processing implementation of the OpenCL standard. It explains the memory structure and gives simple programming examples.

1.1 Terminology

Term	Description
compute kernel	To define a *compute kernel*, it is first necessary to define a kernel. A *kernel* is a small unit of execution that performs a clearly defined function and that can be executed in parallel. Such a kernel can be executed on each element of an input stream (called an NDRange), or simply at each point in an arbitrary index space. A kernel is analogous and, on some devices identical, to what graphics programmers call a shader program. This kernel is not to be confused with an OS kernel, which controls hardware. The most basic form of an NDRange is simply mapped over input data and produces one output item for each input tuple. Subsequent extensions of the basic model provide random-access functionality, variable output counts, and reduction/accumulation operations. Kernels are specified using the kernel keyword. A compute kernel is a specific type of kernel that is not part of the traditional graphics pipeline. The compute kernel type can be used for graphics, but its strength lies in using it for non-graphics fields such as physics, AI, modeling, HPC, and various other computationally intensive applications. In a compute kernel, the work-item spawn order is sequential. This means that on a chip with N work-items per wavefront, the first N work-items go to wavefront 1, the second N work-items go to wavefront 2, etc. Thus, the work-item IDs for wavefront K are in the range (K•N) to ((K+1)•N) - 1.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Term	Description
wavefronts and work-groups	Wavefronts and work-groups are two concepts relating to compute kernels that provide data-parallel granularity. A wavefront executes a number of work-items in lock step relative to each other. Sixteen work-items are executed in parallel across the vector unit, and the whole wavefront is covered over four clock cycles. It is the lowest level that flow control can affect. This means that if two work-items inside of a wavefront go divergent paths of flow control, all work-items in the wavefront go to both paths of flow control. Grouping is a higher-level granularity of data parallelism that is enforced in software, not hardware. Synchronization points in a kernel guarantee that all work-items in a work-group reach that point (barrier) in the code before the next statement is executed. Work-groups are composed of wavefronts. Best performance is attained when the group size is an integer multiple of the wavefront size.
local data store (LDS)	The LDS is a high-speed, low-latency memory private to each compute unit. It is a full gather/scatter model: a work-group can write anywhere in its allocated space. This model is unchanged for the AMD Radeon™ HD 7XXX series. The constraints of the current LDS model are: • The LDS size is allocated per work-group. Each work-group specifies how much of the LDS it requires. The hardware scheduler uses this information to determine which work groups can share a compute unit. • Data can only be shared within work-items in a work-group. • Memory accesses outside of the work-group result in undefined behavior.

1.2 OpenCL Overview

The OpenCL programming model consists of producing complicated task graphs from data-parallel execution nodes.

In a given data-parallel execution, commonly known as a kernel launch, a computation is defined in terms of a sequence of instructions that executes at each point in an N-dimensional index space. It is a common, though by not required, formulation of an algorithm that each computation index maps to an element in an input data set.

The OpenCL data-parallel programming model is hierarchical. The hierarchical subdivision can be specified in two ways:

- Explicitly - the developer defines the total number of work-items to execute in parallel, as well as the division of work-items into specific work-groups.
- Implicitly - the developer specifies the total number of work-items to execute in parallel, and OpenCL manages the division into work-groups.

OpenCL's API also supports the concept of a task dispatch. This is equivalent to executing a kernel on a compute device with a work-group and NDRange containing a single work-item. Parallelism is expressed using vector data types implemented by the device, enqueuing multiple tasks, and/or enqueuing native kernels developed using a programming model orthogonal to OpenCL.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

1.3 Programming Model

The OpenCL programming model is based on the notion of a host device, supported by an application API, and a number of devices connected through a bus. These are programmed using OpenCL C. The host API is divided into platform and runtime layers. OpenCL C is a C-like language with extensions for parallel programming such as memory fence operations and barriers. Figure 1.1 illustrates this model with queues of commands, reading/writing data, and executing kernels for specific devices.

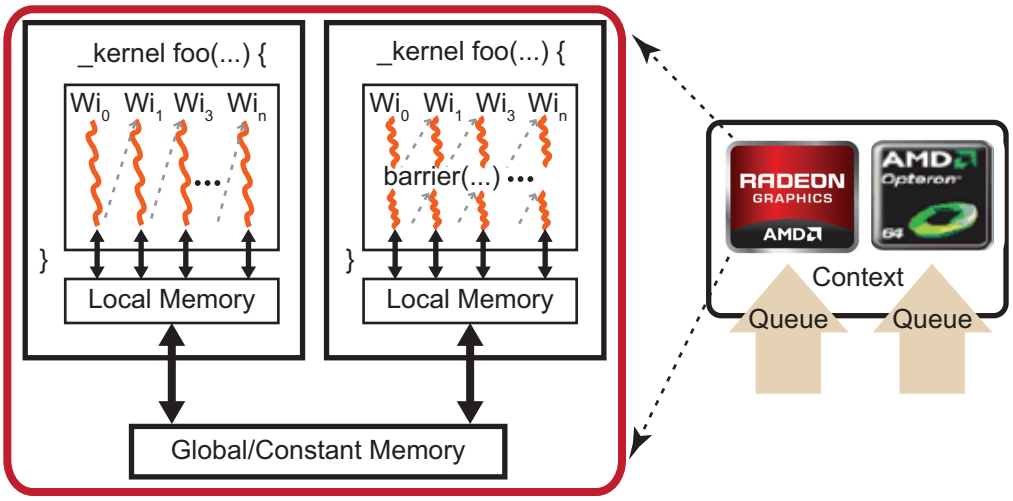

Figure 1.1 OpenCL Programming Model

The devices are capable of running data- and task-parallel work. A kernel can be executed as a function of multi-dimensional domains of indices. Each element is called a *work-item*; the total number of indices is defined as the *global work-size*. The global work-size can be divided into sub-domains, called work-groups, and individual work-items within a group can communicate through global or locally shared memory. Work-items are synchronized through barrier or fence operations. Figure 1.1 is a representation of the host/device architecture with a single platform, consisting of a GPU and a CPU.

An OpenCL application is built by first querying the runtime to determine which platforms are present. There can be any number of different OpenCL implementations installed on a single system. The desired OpenCL platform can be selected by matching the platform vendor string to the desired vendor name, such as “Advanced Micro Devices, Inc.” The next step is to create a context. As shown in Figure 1.1, an OpenCL context has associated with it a number of compute devices (for example, CPU or GPU devices),. Within a context, OpenCL guarantees a relaxed consistency between these devices. This means that memory objects, such as buffers or images, are allocated per context; but changes made by one device are only guaranteed to be visible by another device at well-defined synchronization points. For this, OpenCL provides events, with the ability to synchronize on a given event to enforce the correct order of execution.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Many operations are performed with respect to a given context; there also are many operations that are specific to a device. For example, program compilation and kernel execution are done on a per-device basis. Performing work with a device, such as executing kernels or moving data to and from the device's local memory, is done using a corresponding command queue. A command queue is associated with a single device and a given context; all work for a specific device is done through this interface. Note that while a single command queue can be associated with only a single device, there is no limit to the number of command queues that can point to the same device. For example, it is possible to have one command queue for executing kernels and a command queue for managing data transfers between the host and the device.

Most OpenCL programs follow the same pattern. Given a specific platform, select a device or devices to create a context, allocate memory, create device-specific command queues, and perform data transfers and computations. Generally, the platform is the gateway to accessing specific devices, given these devices and a corresponding context, the application is independent of the platform. Given a context, the application can:

- Create one or more command queues.
- Create programs to run on one or more associated devices.
- Create kernels within those programs.
- Allocate memory buffers or images, either on the host or on the device(s). (Memory can be copied between the host and device.)
- Write data to the device.
- Submit the kernel (with appropriate arguments) to the command queue for execution.
- Read data back to the host from the device.

The relationship between context(s), device(s), buffer(s), program(s), kernel(s), and command queue(s) is best seen by looking at sample code.

1.4 Synchronization

The two domains of synchronization in OpenCL are work-items in a single work-group and command-queue(s) in a single context. Work-group barriers enable synchronization of work-items in a work-group. Each work-item in a work-group must first execute the barrier before executing any instruction beyond this barrier. Either all of, or none of, the work-items in a work-group must encounter the barrier. A barrier or `mem_fence` operation does not have global scope, but is relevant only to the local workgroup on which they operate.

There are two types of synchronization between commands in a command-queue:

- command-queue barrier - enforces ordering within a single queue. Any resulting changes to memory are available to the following commands in the queue.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

- events - enforces ordering between, or within, queues. Enqueued commands in OpenCL return an event identifying the command as well as the memory object updated by it. This ensures that following commands waiting on that event see the updated memory objects before they execute.

1.5 Memory Architecture and Access

OpenCL has four memory domains: private, local, global, and constant; the AMD Accelerated Parallel Processing system also recognizes host (CPU) and PCI Express® (PCIe®) memory.

Memory Type	Description
private	Specific to a work-item; it is not visible to other work-items.
local	Specific to a work-group; accessible only by work-items belonging to that work-group.
global	Accessible to all work-items executing in a context, as well as to the host (read, write, and map commands).
constant	Read-only region for host-allocated and -initialized objects that are not changed during kernel execution.
host (CPU)	Host-accessible region for an application's data structures and program data.
PCIe	Part of host (CPU) memory accessible from, and modifiable by, the host program and the GPU compute device. Modifying this memory requires synchronization between the GPU compute device and the CPU.

Figure 1.2 illustrates the interrelationship of the memories.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

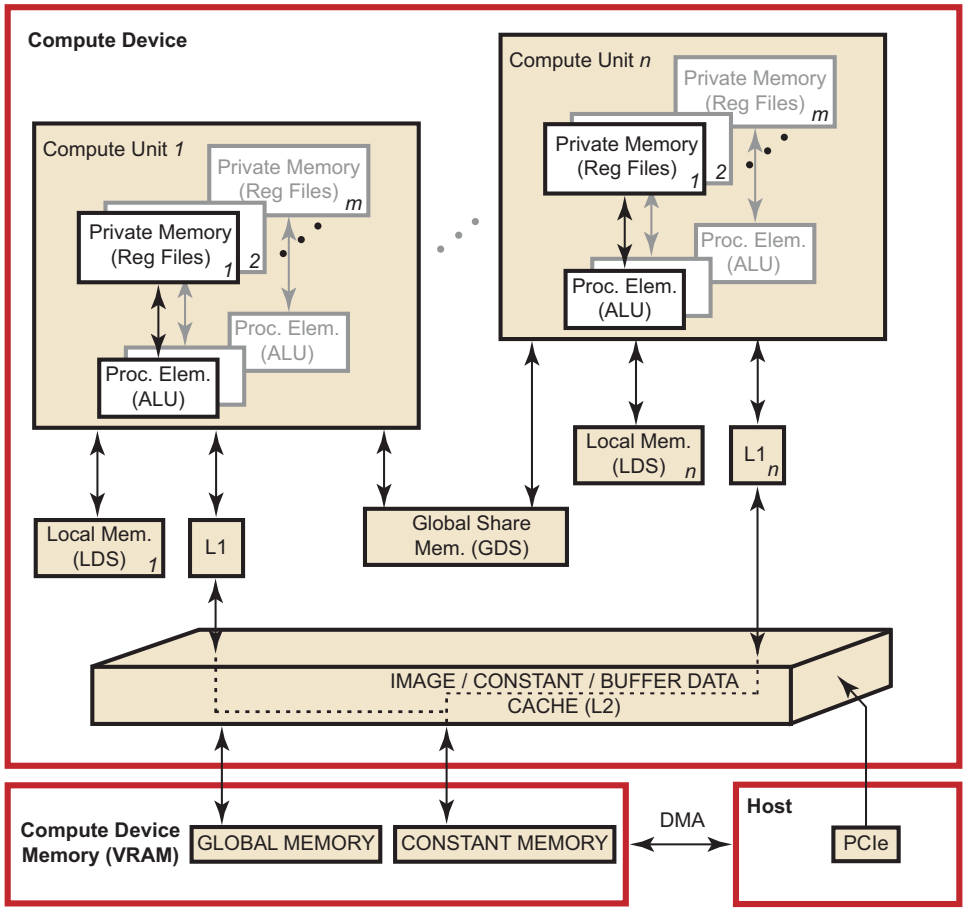

Figure 1.2 Interrelationship of Memory Domains for Southern Islands Devices

Figure 1.3 illustrates the standard dataflow between host (CPU) and GPU.

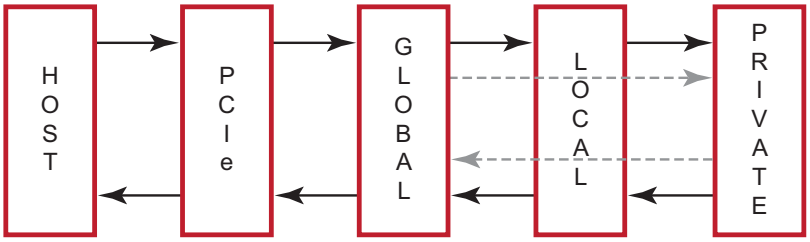

Figure 1.3 Dataflow between Host and GPU

There are two ways to copy data from the host to the GPU compute device memory:

- Implicitly by using `clEnqueueMapBuffer` and `clEnqueueUnMapMemObject`.
- Explicitly through `clEnqueueReadBuffer` and `clEnqueueWriteBuffer` (`clEnqueueReadImage`, `clEnqueueWriteImage`.).

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

When using these interfaces, it is important to consider the amount of copying involved. There is a two-copy processes: between host and PCIe, and between PCIe and GPU compute device.

With proper memory transfer management and the use of system pinned memory (host/CPU memory remapped to the PCIe memory space), copying between host (CPU) memory and PCIe memory can be skipped.

Double copying lowers the overall system memory bandwidth. In GPU compute device programming, pipelining and other techniques help reduce these bottlenecks. See the *AMD OpenCL Optimization Reference Guide* for more specifics about optimization techniques.

1.5.1 Data Share Operations

Local data share (LDS) is a very low-latency, RAM scratchpad for temporary data with at least one order of magnitude higher effective bandwidth than direct, uncached global memory. It permits sharing of data between work-items in a work-group, as well as holding parameters for pixel shader parameter interpolation. Unlike read-only caches, the LDS permits high-speed write-to-read re-use of the memory space (full gather/read/load and scatter/write/store operations).

Figure 1.4 shows the conceptual framework of the LDS is integration into the memory of AMD GPUs using OpenCL.

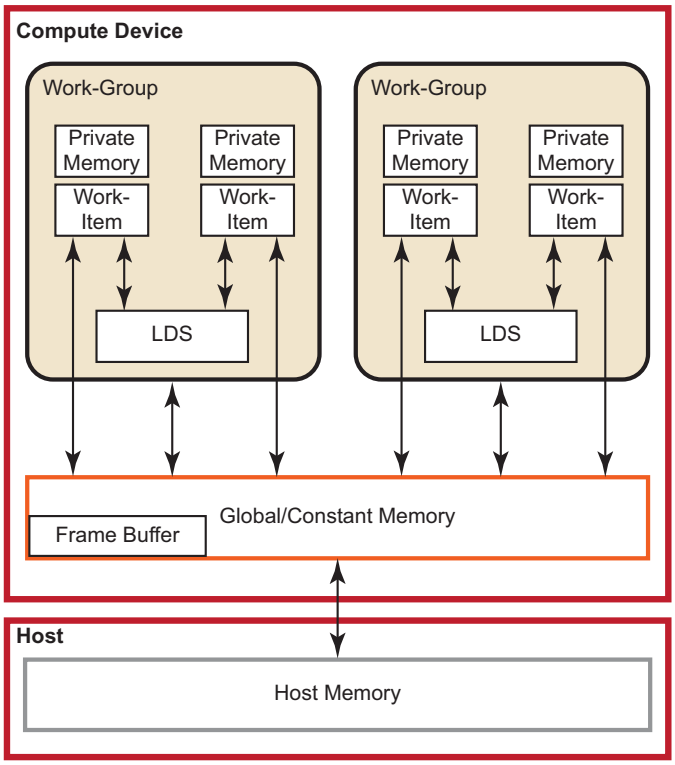

Figure 1.4 High-Level Memory Configuration

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Physically located on-chip, directly next to the ALUs, the LDS is approximately one order of magnitude faster than global memory (assuming no bank conflicts).

There are 32 kB memory per compute unit, segmented into 32 or 16 banks (depending on the GPU type) of 1 k dwords (for 32 banks) or 2 k dwords (for 16 banks). Each bank is a 256x32 two-port RAM (1R/1W per clock cycle). Dwords are placed in the banks serially, but all banks can execute a store or load simultaneously. One work-group can request up to 32 kB memory. Reads across wavefront are dispatched over four cycles in waterfall.

The high bandwidth of the LDS memory is achieved not only through its proximity to the ALUs, but also through simultaneous access to its memory banks. Thus, it is possible to concurrently execute 32 write or read instructions, each nominally 32-bits; extended instructions, read2/write2, can be 64-bits each. If, however, more than one access attempt is made to the same bank at the same time, a bank conflict occurs. In this case, for indexed and atomic operations, hardware prevents the attempted concurrent accesses to the same bank by turning them into serial accesses. This decreases the effective bandwidth of the LDS. For maximum throughput (optimal efficiency), therefore, it is important to avoid bank conflicts. A knowledge of request scheduling and address mapping is key to achieving this.

1.5.2 Dataflow in Memory Hierarchy

Figure 1.5 is a conceptual diagram of the dataflow within the memory structure.

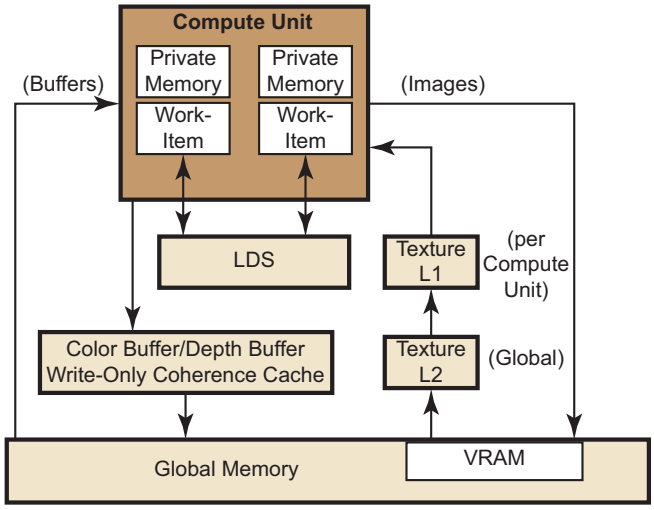

Figure 1.5 Memory Hierarchy Dataflow

To load data into LDS from global memory, it is read from global memory and placed into the work-item's registers; then, a store is performed to LDS. Similarly, to store data into global memory, data is read from LDS and placed into the work-item's registers, then placed into global memory. To make effective use of the LDS, an algorithm must perform many operations on what is transferred between

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

global memory and LDS. It also is possible to load data from a memory buffer directly into LDS, bypassing VGPRs.

LDS atomics are performed in the LDS hardware. (Thus, although ALUs are not directly used for these operations, latency is incurred by the LDS executing this function.) If the algorithm does not require write-to-read reuse (the data is read only), it usually is better to use the image dataflow (see right side of Figure 1.5) because of the cache hierarchy.

Actually, buffer reads may use L1 and L2. When caching is not used for a buffer, reads from that buffer bypass L2. After a buffer read, the line is invalidated; then, on the next read, it is read again (from the same wavefront or from a different clause). After a buffer write, the changed parts of the cache line are written to memory.

Buffers and images are written through the texture L2 cache, but this is flushed immediately after an image write.

The data in private memory is first placed in registers. If more private memory is used than can be placed in registers, or dynamic indexing is used on private arrays, the overflow data is placed (spilled) into scratch memory. Scratch memory is a private subset of global memory, so performance can be dramatically degraded if spilling occurs.

Global memory can be in the high-speed GPU memory (VRAM) or in the host memory, which is accessed by the PCIe bus. A work-item can access global memory either as a buffer or a memory object. Buffer objects are generally read and written directly by the work-items. Data is accessed through the L2 and L1 data caches on the GPU. This limited form of caching provides read coalescing among work-items in a wavefront. Similarly, writes are executed through the texture L2 cache.

Global atomic operations are executed through the texture L2 cache. Atomic instructions that return a value to the kernel are handled similarly to fetch instructions: the kernel must use `S_WAITCNT` to ensure the results have been written to the destination GPR before using the data.

1.5.3 Memory Access

Using local memory (known as local data store, or LDS, as shown in Figure 1.2) typically is an order of magnitude faster than accessing host memory through global memory (VRAM), which is one order of magnitude faster again than PCIe. However, stream cores do not directly access memory; instead, they issue memory requests through dedicated hardware units. When a work-item tries to access memory, the work-item is transferred to the appropriate fetch unit. The work-item then is deactivated until the access unit finishes accessing memory. Meanwhile, other work-items can be active within the compute unit, contributing to better performance. The data fetch units handle three basic types of memory operations: loads, stores, and streaming stores. GPU compute devices can store writes to random memory locations using global buffers.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

1.5.4 Global Memory

The global memory lets applications read from, and write to, arbitrary locations in memory. When using global memory, such read and write operations from the stream kernel are done using regular GPU compute device instructions with the global memory used as the source or destination for the instruction. The programming interface is similar to load/store operations used with CPU programs, where the relative address in the read/write buffer is specified.

When using a global memory, each work-item can write to an arbitrary location within it. Global memory use a linear layout. If consecutive addresses are written, the compute unit issues a burst write for more efficient memory access. Only read-only buffers, such as constants, are cached.

1.5.5 Image Read/Write

Image reads are done by addressing the desired location in the input memory using the fetch unit. The fetch units can process either 1D or 2 D addresses. These addresses can be *normalized* or *un-normalized*. Normalized coordinates are between 0.0 and 1.0 (inclusive). For the fetch units to handle 2D addresses and normalized coordinates, pre-allocated memory segments must be bound to the fetch unit so that the correct memory address can be computed. For a single kernel invocation, up to 128 images can be bound at once for reading, and eight for writing. The maximum number of addresses is 8192x8192 for Evergreen and Northern Islands-based devices, 16384x16384 for SI-based products.

Image reads are cached through the texture system (corresponding to the L2 and L1 caches).

1.6 Example Programs

The following subsections provide simple programming examples with explanatory comments.

1.6.1 First Example: Simple Buffer Write

This sample shows a minimalist OpenCL C program that sets a given buffer to some value. It illustrates the basic programming steps with a minimum amount of code. This sample contains no error checks and the code is not generalized. Yet, many simple test programs might look very similar. The entire code for this sample is provided at the end of this section.

1. The host program must select a platform, which is an abstraction for a given OpenCL implementation. Implementations by multiple vendors can coexist on a host, and the sample uses the first one available.
2. A device id for a GPU device is requested. A CPU device could be requested by using `CL_DEVICE_TYPE_CPU` instead. The device can be a physical device, such as a given GPU, or an abstracted device, such as the collection of all CPU cores on the host.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

3. On the selected device, an OpenCL context is created. A context ties together a device, memory buffers related to that device, OpenCL programs, and command queues. Note that buffers related to a device can reside on either the host or the device. Many OpenCL programs have only a single context, program, and command queue.
4. Before an OpenCL kernel can be launched, its program source is compiled, and a handle to the kernel is created.
5. A memory buffer is allocated in the context.
6. The kernel is launched. While it is necessary to specify the global work size, OpenCL determines a good local work size for this device. Since the kernel was launch asynchronously, `clFinish()` is used to wait for completion.
7. The data is mapped to the host for examination. Calling `clEnqueueMapBuffer` ensures the visibility of the buffer on the host, which in this case probably includes a physical transfer. Alternatively, we could use `clEnqueueWriteBuffer()`, which requires a pre-allocated host-side buffer.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Example Code 1 –

```
//
// Copyright (c) 2010 Advanced Micro Devices, Inc. All rights reserved.
//

// A minimalist OpenCL program.

#include <CL/cl.h>
#include <stdio.h>

#define NWITEMS 512

// A simple memset kernel

const char *source =
"__kernel void memset( __global uint *dst )                                    \n"
"{                                                                             \n"
"   dst[get_global_id(0)] = get_global_id(0);                                  \n"
"}                                                                             \n";

int main(int argc, char ** argv)
{
   // 1. Get a platform.

   cl_platform_id platform;

   clGetPlatformIDs( 1, &platform, NULL );

    // 2. Find a gpu device.

   cl_device_id device;

   clGetDeviceIDs( platform, CL_DEVICE_TYPE_GPU,
                             1,
                             &device,
                             NULL);
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
// Copyright (c) 2010 Advanced Micro Devices, Inc. All rights reserved.
//

// A minimalist OpenCL program.

#include <CL/cl.h>
#include <stdio.h>

#define NWITEMS 512

// A simple memset kernel

const char *source =
"__kernel void memset( __global uint *dst )                                    \n"
"{                                                                             \n"
"   dst[get_global_id(0)] = get_global_id(0);                                  \n"
"}                                                                             \n";

int main(int argc, char ** argv)
{
   // 1. Get a platform.

   cl_platform_id platform;

   clGetPlatformIDs( 1, &platform, NULL );

    // 2. Find a gpu device.

   cl_device_id device;

   clGetDeviceIDs( platform, CL_DEVICE_TYPE_GPU,
                             1,
                             &device,
                             NULL);

   // 3. Create a context and command queue on that device.

   cl_context context = clCreateContext( NULL,
                                         1,
                                         &device,
                                         NULL, NULL, NULL);

   cl_command_queue queue = clCreateCommandQueue( context,
                                                  device,
                                                  0, NULL );

   // 4. Perform runtime source compilation, and obtain kernel entry point.

   cl_program program = clCreateProgramWithSource( context,
                                                   1,
                                                   &source,
                                                   NULL, NULL );

   clBuildProgram( program, 1, &device, NULL, NULL, NULL );

   cl_kernel kernel = clCreateKernel( program, "memset", NULL );

   // 5. Create a data buffer.

   cl_mem buffer = clCreateBuffer( context,
                                   CL_MEM_WRITE_ONLY,
```

//

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
// 6. Launch the kernel. Let OpenCL pick the local work size.

   size_t global_work_size = NWITEMS;

   clSetKernelArg(kernel, 0, sizeof(buffer), (void*) &buffer);

   clEnqueueNDRangeKernel( queue,
                           kernel,
                           1,
                           NULL,
                           &global_work_size,
                           NULL, 0, NULL, NULL);

   clFinish( queue );

   // 7. Look at the results via synchronous buffer map.

   cl_uint *ptr;
   ptr = (cl_uint *) clEnqueueMapBuffer( queue,
                                         buffer,
                                         CL_TRUE,
                                         CL_MAP_READ,
                                         0,
                                         NWITEMS * sizeof(cl_uint),
                                         0, NULL, NULL, NULL );

   int i;

   for(i=0; i < NWITEMS; i++)
      printf("%d %d\n", i, ptr[i]);

   return 0;
}
```

1.6.2 Example: SAXPY Function

This section provides an introductory sample for beginner-level OpenCL programmers using C++ bindings.

The sample implements the SAXPY function (**Y** = *a***X** + **Y**, where **X** and **Y** are vectors, and *a* is a scalar). The full code is reproduced at the end of this section. It uses C++ bindings for OpenCL. These bindings are available in the `CL/cl.hpp` file in the SDK available for AMD Accelerated Parallel Processing; they also are downloadable from the Khronos website: http://www.khronos.org/registry/cl .

The following steps guide you through this example.

1. Enable error checking through the exception handling mechanism in the C++ bindings by using the following define.

   ```
   #define __CL_ENABLE_EXCEPTIONS
   ```

   This removes the need to error check after each OpenCL call. If there is an error, the C++ bindings code throw an exception that is caught at the end of the try block, where we can clean up the host memory allocations. In this example, the C++ objects representing OpenCL resources (`cl::Context`, `cl::CommandQueue`, etc.) are declared as automatic variables, so they do not

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

need to be released. If an OpenCL call returns an error, the error code is defined in the CL/cl.h file.

2. The kernel is very simple: each work-item, *i*, does the SAXPY calculation for its corresponding elements **Y**[i] = *a***X**[i] + **Y**[i]. Both **X** and **Y** vectors are stored in global memory; **X** is read-only, **Y** is read-write.

```
__kernel void saxpy(const __global float * X,
                    __global float * Y,
                    const float a)
{
    uint gid = get_global_id(0);
    Y[gid] = a* X[gid] + Y[gid];
}
```

3. List all platforms on the machine, then select one.

```
cl::Platform::get(&platforms);
```

4. Create an OpenCL context on that platform.

```
cl_context_properties cps[3] = { CL_CONTEXT_PLATFORM,
                                 (cl_context_properties)(*iter)(), 0 };
context = cl::Context(CL_DEVICE_TYPE_GPU, cps);
```

5. Get OpenCL devices from the context.

```
devices = context.getInfo<CL_CONTEXT_DEVICES>();
```

6. Create an OpenCL command queue.

```
queue = cl::CommandQueue(context, devices[0]);
```

7. Create two buffers, corresponding to the **X** and **Y** vectors. Ensure the host-side buffers, pX and pY, are allocated and initialized. The CL_MEM_COPY_HOST_PTR flag instructs the runtime to copy over the contents of the host pointer pX in order to initialize the buffer bufX. The bufX buffer uses the CL_MEM_READ_ONLY flag, while bufY requires the CL_MEM_READ_WRITE flag.

```
bufX = cl::Buffer(context, CL_MEM_READ_ONLY | CL_MEM_COPY_HOST_PTR,
                  sizeof(cl_float) * length, pX);
```

8. Create a program object from the kernel source string, build the program for our devices, and create a kernel object corresponding to the SAXPY kernel. (At this point, it is possible to create multiple kernel objects if there are more than one.)

```
cl::Program::Sources sources(1, std::make_pair(kernelStr.c_str(),
                                kernelStr.length()));
program = cl::Program(context, sources);
program.build(devices);
kernel = cl::Kernel(program, "saxpy");
```

9. Enqueue the kernel for execution on the device (GPU in our example).

   Set each argument individually in separate kernel.setArg() calls. The arguments, do not need to be set again for subsequent kernel enqueue calls. Reset only those arguments that are to pass a new value to the kernel. Then, enqueue the kernel to the command queue with the appropriate global and local work sizes.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
    kernel.setArg(0, bufX);
    kernel.setArg(1, bufY);
    kernel.setArg(2, a);
    queue.enqueueNDRangeKernel(kernel, cl::NDRange(),
                               cl::NDRange(length), cl::NDRange(64));
```

10. Read back the results from bufY to the host pointer pY. We will make this a blocking call (using the `CL_TRUE` argument) since we do not want to proceed before the kernel has finished execution and we have our results back.

```
    queue.enqueueReadBuffer(bufY, CL_TRUE, 0, length * sizeof(cl_float),
pY);
```

11. Clean up the host resources (pX and pY). OpenCL resources is cleaned up by the C++ bindings support code.

The `catch(cl::Error err)` block handles exceptions thrown by the C++ bindings code. If there is an OpenCL call error, it prints out the name of the call and the error code (codes are defined in `CL/cl.h`). If there is a kernel compilation error, the error code is `CL_BUILD_PROGRAM_FAILURE`, in which case it is necessary to print out the build log.

Example Code 2 –

```
#define __CL_ENABLE_EXCEPTIONS

#include <CL/cl.hpp>
#include <string>
#include <iostream>
#include <string>

using std::cout;
using std::cerr;
using std::endl;
using std::string;

///////////////////////////////////////////////////////////////
// Helper function to print vector elements
///////////////////////////////////////////////////////////////
void printVector(const std::string arrayName,
                 const cl_float * arrayData,
                 const unsigned int length)
{
    int numElementsToPrint = (256 < length) ? 256 : length;
    cout << endl << arrayName << ":" << endl;
    for(int i = 0; i < numElementsToPrint; ++i)
        cout << arrayData[i] << " ";
    cout << endl;
}

///////////////////////////////////////////////////////////////
// Globals
///////////////////////////////////////////////////////////////
int length        = 256;
cl_float * pX     = NULL;
cl_float * pY     = NULL;
cl_float a        = 2.f;



std::vector<cl::Platform> platforms;
cl::Context        context;
std::vector<cl::Device> devices;
cl::CommandQueue   queue;
cl::Program        program;
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
cl::Kernel         kernel;
cl::Buffer         bufX;
cl::Buffer         bufY;

///////////////////////////////////////////////////////////////////
// The saxpy kernel
///////////////////////////////////////////////////////////////////
string kernelStr    =
    "__kernel void saxpy(const __global float * x,\n"
    "                    __global float * y,\n"
    "                    const float a)\n"
    "{\n"
    "    uint gid = get_global_id(0);\n"
    "    y[gid] = a* x[gid] + y[gid];\n"
    "}\n";

///////////////////////////////////////////////////////////////////
// Allocate and initialize memory on the host
///////////////////////////////////////////////////////////////////
void initHost()
{
    size_t sizeInBytes = length * sizeof(cl_float);
    pX = (cl_float *) malloc(sizeInBytes);
    if(pX == NULL)
        throw(string("Error: Failed to allocate input memory on host\n"));

    pY = (cl_float *) malloc(sizeInBytes);
    if(pY == NULL)
        throw(string("Error: Failed to allocate input memory on host\n"));

    for(int i = 0; i < length; i++)
    {
        pX[i] = cl_float(i);
        pY[i] = cl_float(length-1-i);
    }

    printVector("X", pX, length);
    printVector("Y", pY, length);
}

///////////////////////////////////////////////////////////////////
// Release host memory
///////////////////////////////////////////////////////////////////
void cleanupHost()
{
    if(pX)
    {
        free(pX);
        pX = NULL;
    }
    if(pY != NULL)
    {
        free(pY);
        pY = NULL;
    }
}


void
main(int argc, char * argv[])




{
    try
    {
        ///////////////////////////////////////////////////////////////
        // Allocate and initialize memory on the host
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
        ///////////////////////////////////////////////////////////////////
        initHost();

        ///////////////////////////////////////////////////////////////////
        // Find the platform
        ///////////////////////////////////////////////////////////////////
        cl::Platform::get(&platforms);
        std::vector<cl::Platform>::iterator iter;
        for(iter = platforms.begin(); iter != platforms.end(); ++iter)
        {
           if(!strcmp((*iter).getInfo<CL_PLATFORM_VENDOR>().c_str(),
           "Advanced Micro Devices, Inc."))
        {
break;
} }

        ///////////////////////////////////////////////////////////////////
        // Create an OpenCL context
        ///////////////////////////////////////////////////////////////////
        cl_context_properties cps[3] = { CL_CONTEXT_PLATFORM,
(cl_context_properties)(*iter)(), 0 };
        context = cl::Context(CL_DEVICE_TYPE_GPU, cps);

        ///////////////////////////////////////////////////////////////////
        // Detect OpenCL devices
        ///////////////////////////////////////////////////////////////////
        devices = context.getInfo<CL_CONTEXT_DEVICES>();

        ///////////////////////////////////////////////////////////////////
        // Create an OpenCL command queue
        ///////////////////////////////////////////////////////////////////
        queue = cl::CommandQueue(context, devices[0]);

        ///////////////////////////////////////////////////////////////////
        // Create OpenCL memory buffers
        ///////////////////////////////////////////////////////////////////
        bufX = cl::Buffer(context,
                          CL_MEM_READ_ONLY | CL_MEM_COPY_HOST_PTR,
                          sizeof(cl_float) * length,
                          pX);
        bufY = cl::Buffer(context,
                          CL_MEM_READ_WRITE | CL_MEM_COPY_HOST_PTR,
                          sizeof(cl_float) * length,
                          pY);

        ///////////////////////////////////////////////////////////////////
        // Load CL file, build CL program object, create CL kernel object
        ///////////////////////////////////////////////////////////////////
        cl::Program::Sources sources(1, std::make_pair(kernelStr.c_str(),
                                       kernelStr.length()));
        program = cl::Program(context, sources);
        program.build(devices);
        kernel = cl::Kernel(program, "saxpy");

        ///////////////////////////////////////////////////////////////////
        // Set the arguments that will be used for kernel execution
        ///////////////////////////////////////////////////////////////////
        kernel.setArg(0, bufX);
        kernel.setArg(1, bufY);
        kernel.setArg(2, a);

        ///////////////////////////////////////////////////////////////////
        // Enqueue the kernel to the queue
        // with appropriate global and local work sizes
        ///////////////////////////////////////////////////////////////////
        queue.enqueueNDRangeKernel(kernel, cl::NDRange(),
                                   cl::NDRange(length), cl::NDRange(64));

        ///////////////////////////////////////////////////////////////////
        // Enqueue blocking call to read back buffer Y
        ///////////////////////////////////////////////////////////////////
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
        queue.enqueueReadBuffer(bufY, CL_TRUE, 0, length *
                                sizeof(cl_float), pY);

        printVector("Y", pY, length);

        ///////////////////////////////////////////////////////////////
        // Release host resources
        ///////////////////////////////////////////////////////////////
        cleanupHost();
    }
    catch (cl::Error err)
    {
        ///////////////////////////////////////////////////////////////
        // Catch OpenCL errors and print log if it is a build error
        ///////////////////////////////////////////////////////////////
        cerr << "ERROR: " << err.what() << "(" << err.err() << ")" <<
            endl;
        if (err.err() == CL_BUILD_PROGRAM_FAILURE)
        {
            string str =
               program.getBuildInfo<CL_PROGRAM_BUILD_LOG>(devices[0]);
            cout << "Program Info: " << str << endl;
        }
        cleanupHost();
    }
    catch(string msg)
    {
        cerr << "Exception caught in main(): " << msg << endl;
        cleanupHost();
    }
}
```

1.6.3 Example: Parallel Min() Function

This medium-complexity sample shows how to implement an efficient parallel min() function.

The code is written so that it performs very well on either CPU or GPU. The number of threads launched depends on how many hardware processors are available. Each thread walks the source buffer, using a device-optimal access pattern selected at runtime. A multi-stage reduction using `__local` and `__global` atomics produces the single result value.

The sample includes a number of programming techniques useful for simple tests. Only minimal error checking and resource tear-down is used.

Runtime Code –

1. The source memory buffer is allocated, and initialized with a random pattern. Also, the actual min() value for this data set is serially computed, in order to later verify the parallel result.
2. The compiler is instructed to dump the intermediate IL and ISA files for further analysis.
3. The main section of the code, including device setup, CL data buffer creation, and code compilation, is executed for each device, in this case for CPU and GPU. Since the source memory buffer exists on the host, it is shared. All other resources are device-specific.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

4. The global work size is computed for each device. A simple heuristic is used to ensure an optimal number of threads on each device. For the CPU, a given CL implementation can translate one work-item per CL compute unit into one thread per CPU core.

   On the GPU, an initial multiple of the wavefront size is used, which is adjusted to ensure even divisibility of the input data over all threads. The value of 7 is a minimum value to keep all independent hardware units of the compute units busy, and to provide a minimum amount of memory latency hiding for a kernel with little ALU activity.

5. After the kernels are built, the code prints errors that occurred during kernel compilation and linking.

6. The main loop is set up so that the measured timing reflects the actual kernel performance. If a sufficiently large NLOOPS is chosen, effects from kernel launch time and delayed buffer copies to the device by the CL runtime are minimized. Note that while only a single `clFinish()` is executed at the end of the timing run, the two kernels are always linked using an *event* to ensure serial execution.

   The bandwidth is expressed as "number of input bytes processed." For high-end graphics cards, the bandwidth of this algorithm is about an order of magnitude higher than that of the CPU, due to the parallelized memory subsystem of the graphics card.

7. The results then are checked against the comparison value. This also establishes that the result is the same on both CPU and GPU, which can serve as the first verification test for newly written kernel code.

8. Note the use of the debug buffer to obtain some runtime variables. Debug buffers also can be used to create short execution traces for each thread, assuming the device has enough memory.

9. You can use the `Timer.cpp` and `Timer.h` files from the TransferOverlap sample, which is in the SDK samples.

Kernel Code –

10. The code uses four-component vectors (uint4) so the compiler can identify concurrent execution paths as often as possible. On the GPU, this can be used to further optimize memory accesses and distribution across ALUs. On the CPU, it can be used to enable SSE-like execution.

11. The kernel sets up a memory access pattern based on the device. For the CPU, the source buffer is chopped into continuous buffers: one per thread. Each CPU thread serially walks through its buffer portion, which results in good cache and prefetch behavior for each core.

    On the GPU, each thread walks the source buffer using a stride of the total number of threads. As many threads are executed in parallel, the result is a maximally coalesced memory pattern requested from the memory back-end. For example, if each compute unit has 16 physical processors, 16 uint4 requests are produced in parallel, per clock, for a total of 256 bytes per clock.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

12. The kernel code uses a reduction consisting of three stages: __global to __private, __private to _local, which is flushed to __global, and finally __global to __global. In the first loop, each thread walks __global memory, and reduces all values into a min value in __private memory (typically, a register). This is the bulk of the work, and is mainly bound by __global memory bandwidth. The subsequent reduction stages are brief in comparison.
13. Next, all per-thread minimum values inside the work-group are reduced to a __local value, using an atomic operation. Access to the __local value is serialized; however, the number of these operations is very small compared to the work of the previous reduction stage. The threads within a work-group are synchronized through a local barrier(). The reduced min value is stored in __global memory.
14. After all work-groups are finished, a second kernel reduces all work-group values into a single value in __global memory, using an atomic operation. This is a minor contributor to the overall runtime.

Example Code 3 –

```
//
// Copyright (c) 2010 Advanced Micro Devices, Inc. All rights reserved.
//

#include <CL/cl.h>
#include <stdio.h>
#include <stdlib.h>
#include <time.h>
#include "Timer.h"

#define NDEVS       2

// A parallel min() kernel that works well on CPU and GPU

const char *kernel_source =
"                                                                               \n"
"#pragma OPENCL EXTENSION cl_khr_local_int32_extended_atomics : enable      \n"
"#pragma OPENCL EXTENSION cl_khr_global_int32_extended_atomics : enable     \n"
"                                                                               \n"
"  // 9. The source buffer is accessed as 4-vectors.                          \n"
"                                                                               \n"
"__kernel void minp( __global uint4 *src,                                    \n"
"                    __global uint  *gmin,                                   \n"
"                    __local  uint  *lmin,                                   \n"
"                    __global uint  *dbg,                                    \n"
"                    int           nitems,                                  \n"
"                    uint           dev )                                    \n"
"{                                                                              \n"
"   // 10. Set up __global memory access pattern.                             \n"
"                                                                               \n"
"   uint  count  = ( nitems / 4 ) / get_global_size(0);                      \n"
"   uint  idx    = (dev == 0) ? get_global_id(0) * count                     \n"
"                             : get_global_id(0);                            \n"
"   uint  stride = (dev == 0) ? 1 : get_global_size(0);                      \n"
"   uint  pmin   = (uint) -1;                                                \n"
"                                                                               \n"
"   // 11. First, compute private min, for this work-item.                    \n"
"                                                                               \n"
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
"   for( int n=0; n < count; n++, idx += stride )                          \n"
"   {                                                                      \n"
"      pmin = min( pmin, src[idx].x );                                     \n"
"      pmin = min( pmin, src[idx].y );                                     \n"
"      pmin = min( pmin, src[idx].z );                                     \n"
"      pmin = min( pmin, src[idx].w );                                     \n"
"   }                                                                      \n"
"                                                                          \n"
"   // 12. Reduce min values inside work-group.                            \n"
"                                                                          \n"
"   if( get_local_id(0) == 0 )                                             \n"
"      lmin[0] = (uint) -1;                                                \n"
"                                                                          \n"
"   barrier( CLK_LOCAL_MEM_FENCE );                                        \n"
"                                                                          \n"
"   (void) atom_min( lmin, pmin );                                         \n"
"                                                                          \n"
"   barrier( CLK_LOCAL_MEM_FENCE );                                        \n"
"                                                                          \n"
"   // Write out to __global.                                              \n"
"                                                                          \n"
"   if( get_local_id(0) == 0 )                                             \n"
"      gmin[ get_group_id(0) ] = lmin[0];                                  \n"
"                                                                          \n"
"   // Dump some debug information.                                        \n"
"                                                                          \n"
"   if( get_global_id(0) == 0 )                                            \n"
"   {                                                                      \n"
"      dbg[0] = get_num_groups(0);                                         \n"
"      dbg[1] = get_global_size(0);                                        \n"
"      dbg[2] = count;                                                     \n"
"      dbg[3] = stride;                                                    \n"
"   }                                                                      \n"
"}                                                                         \n"
"                                                                          \n"
"// 13. Reduce work-group min values from __global to __global.           \n"
"                                                                          \n"
"__kernel void reduce( __global uint4 *src,                                \n"
"                      __global uint  *gmin )                              \n"
"{                                                                         \n"
"   (void) atom_min( gmin, gmin[get_global_id(0)] ) ;                      \n"
"}                                                                        \n";


int main(int argc, char ** argv)
{
   cl_platform_id   platform;

   int              dev, nw;
   cl_device_type   devs[NDEVS] = { CL_DEVICE_TYPE_CPU,
                                    CL_DEVICE_TYPE_GPU };

   cl_uint          *src_ptr;
   unsigned int     num_src_items = 4096*4096;

   // 1. quick & dirty MWC random init of source buffer.

   // Random seed (portable).

   time_t ltime;
   time(&ltime);

   src_ptr = (cl_uint *) malloc( num_src_items * sizeof(cl_uint) );
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
   cl_uint a =   (cl_uint) ltime,
           b =   (cl_uint) ltime;
   cl_uint min = (cl_uint) -1;

   // Do serial computation of min() for result verification.

   for( int i=0; i < num_src_items; i++ )
   {
      src_ptr[i] = (cl_uint) (b = ( a * ( b & 65535 )) + (  b >> 16 ));
      min = src_ptr[i] < min ? src_ptr[i] : min;
   }

// Get a platform.

   clGetPlatformIDs( 1, &platform, NULL );

   // 3. Iterate over devices.

   for(dev=0; dev < NDEVS; dev++)
   {
      cl_device_id     device;
      cl_context       context;
      cl_command_queue queue;
      cl_program       program;
      cl_kernel        minp;
      cl_kernel        reduce;

      cl_mem           src_buf;
      cl_mem           dst_buf;
      cl_mem           dbg_buf;

      cl_uint          *dst_ptr,
                       *dbg_ptr;

      printf("\n%s: ", dev == 0 ? "CPU" : "GPU");

      // Find the device.

      clGetDeviceIDs( platform,
                      devs[dev],
                      1,
                      &device,
                      NULL);

      // 4. Compute work sizes.

      cl_uint compute_units;
      size_t  global_work_size;
      size_t  local_work_size;
      size_t  num_groups;

      clGetDeviceInfo( device,
                       CL_DEVICE_MAX_COMPUTE_UNITS,
                       sizeof(cl_uint),
                       &compute_units,
                       NULL);

      if( devs[dev] == CL_DEVICE_TYPE_CPU )
      {
         global_work_size = compute_units * 1;      // 1 thread per core
         local_work_size = 1;
      }
      else
      {
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
    cl_uint ws = 64;

    global_work_size = compute_units * 7 * ws; // 7 wavefronts per SIMD

    while( (num_src_items / 4) % global_work_size != 0 )
       global_work_size += ws;

    local_work_size = ws;
}

num_groups = global_work_size / local_work_size;

// Create a context and command queue on that device.

context = clCreateContext( NULL,
                           1,
                           &device,
                           NULL, NULL, NULL);

queue = clCreateCommandQueue(context,
                             device,
                             0, NULL);

// Minimal error check.

if( queue == NULL )
{
    printf("Compute device setup failed\n");
    return(-1);
}

// Perform runtime source compilation, and obtain kernel entry point.

program = clCreateProgramWithSource( context,
                                     1,
                                     &kernel_source,
                                     NULL, NULL );

//Tell compiler to dump intermediate .il and .isa GPU files.

ret = clBuildProgram( program,
                      1,
                      &device,
                      "-save-temps",
                      NUL, NULL );

// 5. Print compiler error messages

if(ret != CL_SUCCESS)
{
   printf("clBuildProgram failed: %d\n", ret);

   char buf[0x10000];

   clGetProgramBuildInfo( program,
                          device,
                          CL_PROGRAM_BUILD_LOG,
                          0x10000,
                          buf,
                          NULL);
   printf("\n%s\n", buf);
   return(-1);
}
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
minp   = clCreateKernel( program, "minp", NULL );
reduce = clCreateKernel( program, "reduce", NULL );

// Create input, output and debug buffers.

src_buf = clCreateBuffer( context,
                          CL_MEM_READ_ONLY | CL_MEM_COPY_HOST_PTR,
                          num_src_items * sizeof(cl_uint),
                          src_ptr,
                          NULL );

dst_buf = clCreateBuffer( context,
                          CL_MEM_READ_WRITE,
                          num_groups * sizeof(cl_uint),
                          NULL, NULL );

dbg_buf = clCreateBuffer( context,
                          CL_MEM_WRITE_ONLY,
                          global_work_size * sizeof(cl_uint),
                          NULL, NULL );

clSetKernelArg(minp, 0, sizeof(void *),        (void*) &src_buf);
clSetKernelArg(minp, 1, sizeof(void *),        (void*) &dst_buf);
clSetKernelArg(minp, 2, 1*sizeof(cl_uint),     (void*) NULL);
clSetKernelArg(minp, 3, sizeof(void *),        (void*) &dbg_buf);
clSetKernelArg(minp, 4, sizeof(num_src_items), (void*) &num_src_items);
clSetKernelArg(minp, 5, sizeof(dev),           (void*) &dev);

clSetKernelArg(reduce, 0, sizeof(void *),      (void*) &src_buf);
clSetKernelArg(reduce, 1, sizeof(void *),      (void*) &dst_buf);

CPerfCounter t;
t.Reset();
t.Start();

// 6. Main timing loop.

#define NLOOPS 500

cl_event ev;
int nloops = NLOOPS;

while(nloops--)
{

   clEnqueueNDRangeKernel( queue,
                           minp,
                           1,
                           NULL,
                           &global_work_size,
                           &local_work_size,
                           0, NULL, &ev);

   clEnqueueNDRangeKernel( queue,
                           reduce,
                           1,
                           NULL,
                           &num_groups,
                           NULL, 1, &ev, NULL);
}

clFinish( queue );
t.Stop();
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
      printf("B/W %.2f GB/sec, ", ((float) num_src_items *
                                   sizeof(cl_uint) * NLOOPS) /
                                  t.GetElapsedTime() / 1e9 );

      // 7. Look at the results via synchronous buffer map.

      dst_ptr = (cl_uint *) clEnqueueMapBuffer( queue,
                                                dst_buf,
                                                CL_TRUE,
                                                CL_MAP_READ,
                                                0,
                                                num_groups * sizeof(cl_uint),
                                                0, NULL, NULL, NULL );

      dbg_ptr = (cl_uint *) clEnqueueMapBuffer( queue,
                                                dbg_buf,
                                                CL_TRUE,
                                                CL_MAP_READ,
                                                0,
                                                global_work_size *
                                                sizeof(cl_uint),
                                                0, NULL, NULL, NULL );

      // 8. Print some debug info.

      printf("%d groups, %d threads, count %d, stride %d\n", dbg_ptr[0],
                                                             dbg_ptr[1],
                                                             dbg_ptr[2],
                                                             dbg_ptr[3] );

      if( dst_ptr[0] == min )
         printf("result correct\n");
      else
         printf("result INcorrect\n");

   }

   printf("\n");
   return 0;
}
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Chapter 2
AMD Implementation

2.1 The AMD Accelerated Parallel Processing Implementation of OpenCL

AMD Accelerated Parallel Processing harnesses the tremendous processing power of GPUs for high-performance, data-parallel computing in a wide range of applications. The AMD Accelerated Parallel Processing system includes a software stack, AMD GPUs, and AMD multicore CPUs. Figure 2.1 illustrates the relationship of the AMD Accelerated Parallel Processing components.

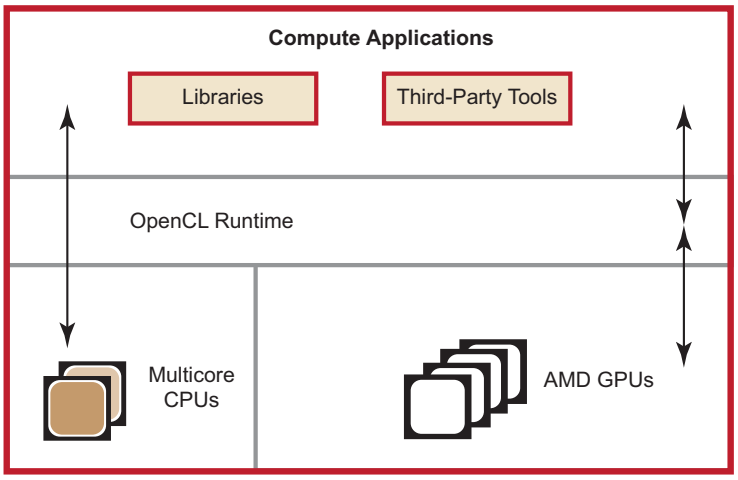

Figure 2.1 AMD Accelerated Parallel Processing Software Ecosystem

The AMD Accelerated Parallel Processing software stack provides end-users and developers with a complete, flexible suite of tools to leverage the processing power in AMD GPUs. AMD Accelerated Parallel Processing software embraces open-systems, open-platform standards. The AMD Accelerated Parallel Processing open platform strategy enables AMD technology partners to develop and provide third-party development tools.

The software includes the following components:

- OpenCL compiler and runtime
- Debugging and Performance Profiling Tools – AMD CodeXL.
- Performance Libraries – AMD Accelerated Parallel Processing Math Library (APPML) for optimized NDRange-specific algorithms.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

The latest generations of AMD GPUs use unified shader architectures capable of running different kernel types interleaved on the same hardware. Programmable GPU compute devices execute various user-developed programs, known to graphics programmers as *shaders* and to compute programmers as *kernels*. These GPU compute devices can execute non-graphics functions using a data-parallel programming model that maps executions onto compute units. Each compute unit contains one (pre-GCN devices) or more (GCN devices) vector (SIMD) units. In this programming model, known as AMD Accelerated Parallel Processing, arrays of input data elements stored in memory are accessed by a number of compute units.

Each instance of a kernel running on a compute unit is called a work-item. Work-items are mapped to an n-dimensional index space, called an *NDRange*.

The GPU schedules the range of work-items onto a group of processing elements, until all work-items have been processed. Subsequent kernels then can be executed, until the application completes. A simplified view of the AMD Accelerated Parallel Processing programming model and the mapping of work-items to processing elements is shown in Figure 2.2.

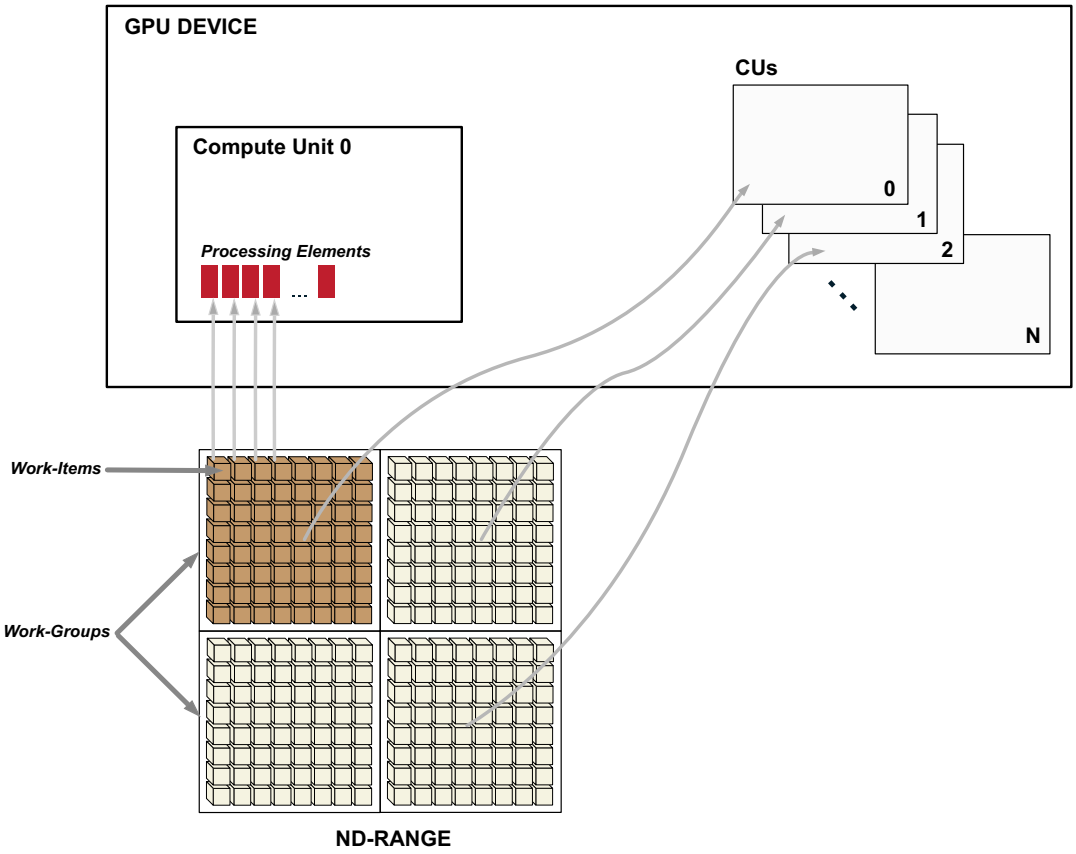

Figure 2.2 Simplified Mapping of OpenCL onto AMD Accelerated Parallel Processing

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Work-groups are assigned to CUs. All work-items of a work-group can be processed only by the processing elements of a single CU. A processing element can process only one work-item at a time; however, a CU can process multiple work-groups.

Note that in OpenCL 2.0, the work-groups are not required to divide evenly into the NDRange.

OpenCL maps the total number of work-items to be launched onto an n-dimensional grid (ND-Range). The developer can specify how to divide these items into work-groups. AMD GPUs execute on wavefronts (groups of work-items executed in lock-step in a compute unit); there is an integer number of wavefronts in each work-group. Thus, as shown in Figure 2.3, hardware that schedules work-items for execution in the AMD Accelerated Parallel Processing environment includes the intermediate step of specifying wavefronts within a work-group. This permits achieving maximum performance on AMD GPUs. For a more detailed discussion of wavefronts, see Section 1.1, “Terminology,” page 1-1.

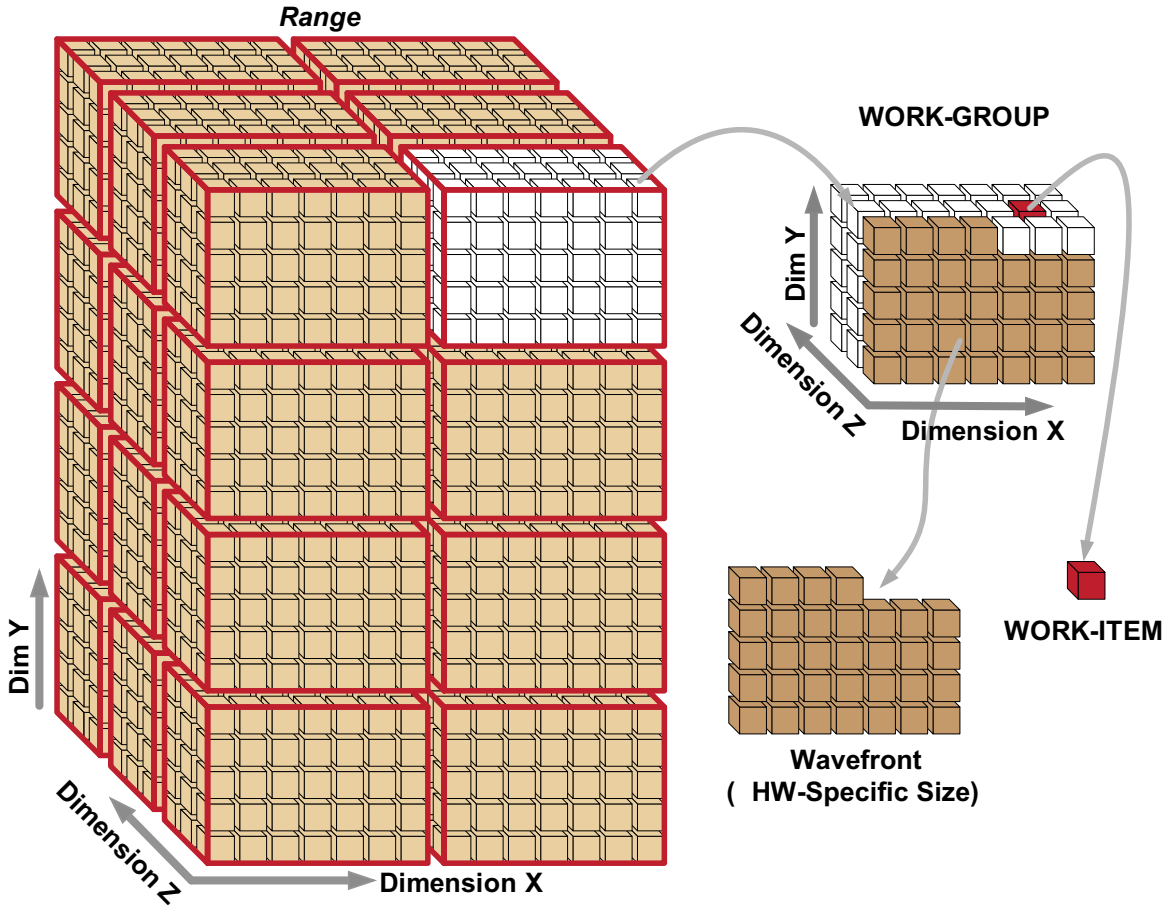

Figure 2.3 Work-Item Grouping Into Work-Groups and Wavefronts

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

2.1.1 Work-Item Processing

All processing elements within a vector unit execute the same instruction in each cycle. A work item can issue one instruction per clock cycle. The block of work-items that are executed together is called a *wavefront*. To hide latencies due to memory accesses and processing element operations, up to four work-items from the same wavefront are pipelined on the same vector unit. For example, on the AMD Radeon™ HD 290X compute device, the 16 processing elements within each vector unit execute the same instruction for four cycles, which effectively appears as a 64-wide compute unit in execution width.

The size of wavefronts can differ on different GPU compute devices. For example, some of the low-end and older GPUs, such as the AMD Radeon™ HD 54XX series graphics cards, have a wavefront size of 32 work-items. Higher-end and newer AMD GPUs have a wavefront size of 64 work-items.

Compute units operate independently of each other, so it is possible for different compute units to execute different instructions. It is also possible for different vector units within a compute unit to execute different instructions.

Before discussing flow control, it is necessary to clarify the relationship of a wavefront to a work-group. If a user defines a work-group, it consists of one or more wavefronts. A wavefront is a hardware thread with its own program counter; it is capable of following control flow independently of other wavefronts. A wavefront consists of 64 or fewer work-items. The mapping is based on a linear work-item order. On a device with a wavefront size of 64, work-items 0-63 map to wavefront 0, work items 64-127 map to wavefront 1, etc. For optimum hardware usage, an integer multiple of 64 work-items is recommended.

2.1.2 Work-Item Creation

For each work-group, the GPU compute device spawns the required number of wavefronts on a single compute unit. If there are non-active work-items within a wavefront, the processing elemnts that would have been mapped to those work-items are idle. An example is a work-group that is a non-multiple of a wavefront size (for example: if the work-group size is 32, the wavefront is half empty and unused).

2.1.3 Flow Control

Flow control, such as branching, is achieved by combining all necessary paths as a wavefront. If work-items within a wavefront diverge, all paths are executed serially. For example, if a work-item contains a branch with two paths, the wavefront first executes one path, then the second path. The total time to execute the branch is the sum of each path time. An important point is that even if only one work-item in a wavefront diverges, the rest of the work-items in the wavefront execute the branch. The number of work-items that must be executed during a branch is called the *branch granularity*. On AMD hardware, the branch granularity is the same as the number of work-items in a wavefront.

Masking of wavefronts is effected by constructs such as:

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
if(x)
  {
   .   //items within these braces = A
   .
   .
  }
else
  {
   .   //items within these braces = B
   .
   .
  }
```

The wavefront mask is set true for lanes (elements/items) in which `x` is true, then execute A. The mask then is inverted, and B is executed.

Example 1: If two branches, A and B, take the same amount of time *t* to execute over a wavefront, the total time of execution, if any work-item diverges, is *2t*.

Loops execute in a similar fashion, where the wavefront occupies a compute unit as long as there is at least one work-item in the wavefront still being processed. Thus, the total execution time for the wavefront is determined by the work-item with the longest execution time.

Example 2: If *t* is the time it takes to execute a single iteration of a loop; and within a wavefront all work-items execute the loop one time, except for a single work-item that executes the loop 100 times, the time it takes to execute that entire wavefront is *100t*.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

2.2 Hardware Overview for GCN Devices

A general OpenCL device comprises compute units (CUs), each of which has sub-modules that ultimately have ALUs. A work-item (or SPMD kernel instance) executes on an ALU, as shown in Figure 2.4).

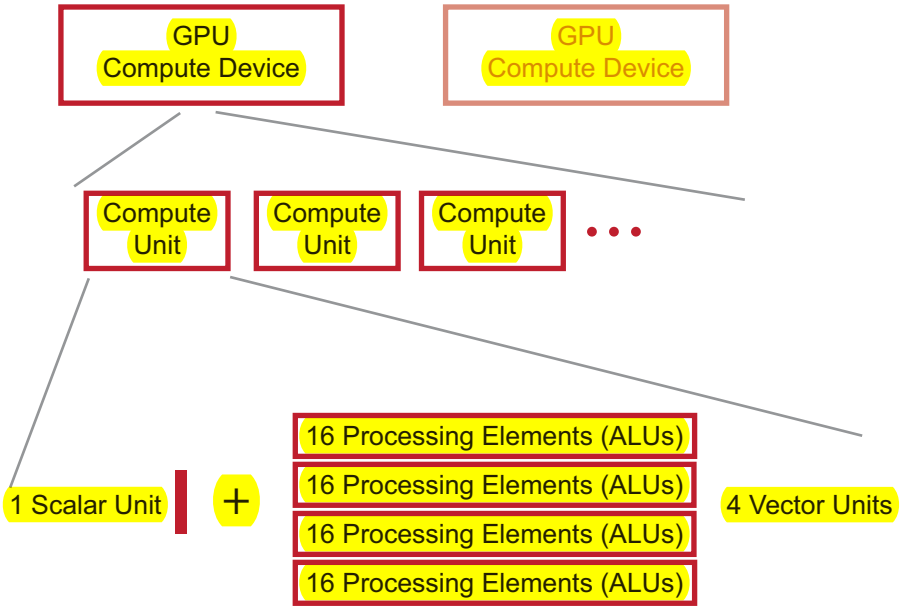

Figure 2.4 Generalized AMD GPU Compute Device Structure for GCN Devices

In GCN devices, each CU includes one Scalar Unit and four Vector (SIMD) units, each of which contains an array of 16 processing elements (PEs). Each PE contains one ALU. Each SIMD unit simultaneously executes a single operation across 16 work items, but each can be working on a separate wavefront.

For example, for the AMD Radeon™ HD 79XX devices each of the 32 CUs has one Scalar Unit and four Vector Units. Figure 2.5 shows only two compute engines/command processors of the array that comprises the compute device of the AMD Radeon™ HD 79XX family.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

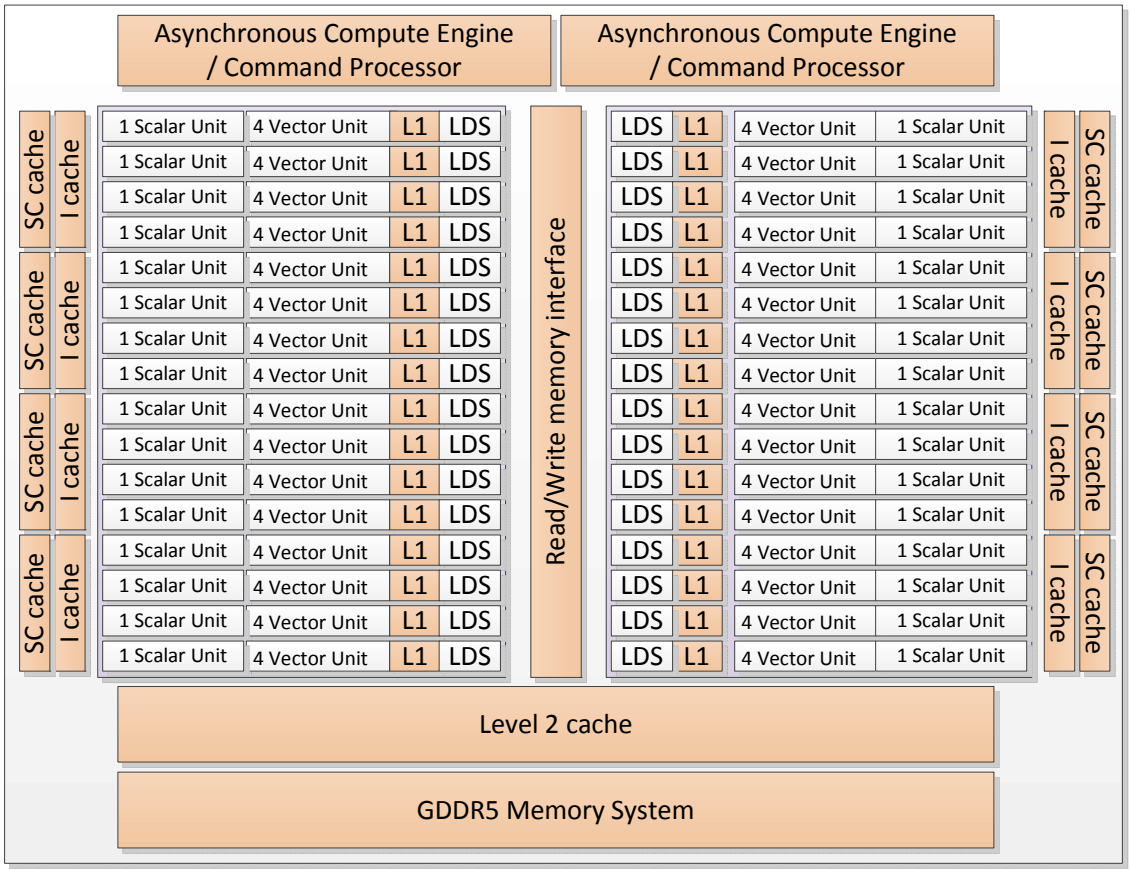

Figure 2.5 AMD Radeon™ HD 79XX Device Partial Block Diagram

In Figure 2.5, there are two command processors, which can process two command queues concurrently. The Scalar Unit, Vector Unit, Level 1 data cache (L1), and Local Data Share (LDS) are the components of one compute unit, of which there are 32. The scalar (SC) cache is the scalar unit data cache, and the Level 2 cache consists of instructions and data.

On GCN devices, the instruction stream contains both scalar and vector instructions. On each cycle, it selects a scalar instruction and a vector instruction (as well as a memory operation and a branch operation, if available); it issues one to the scalar unit, the other to the vector unit; this takes four cycles to issue over the four vector cores (the same four cycles over which the 16 units execute 64 work-items).

The Asynchronous Compute Engines (ACEs) manage the CUs; a graphics command processor handles graphics shaders and fixed-function hardware.

2.2.1 Key differences between pre-GCN and GCN devices

In pre-GCN devices (for a hardware overview, see Appendix D, “Hardware overview of pre-GCN devices.”), each compute unit consists of a single vector unit, each containing up to 16 processing elements. Each processing element,

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

which contains 4 or 5 ALUs, could execute bundles of 4 or 5 independent instructions co-issued in a VLIW (Very Long Instruction Word) format. All the processing elements within a vector unit execute a single wavefront (a group of 64 work items). If operations within a wavefront contain dependencies, they cannot be scheduled in the same clock cycle, leaving some ALUs un-utilized. In such cases, some processing elements (and hence, vector units) remain under-utilized.

In GCN devices, the CUs are arranged in four vector unit arrays consisting of 16 processing elements each. Each of these arrays executes a single instruction across each lane for each block of 16 work-items. That instruction is repeated over four cycles to make the 64-element vector called a *wavefront*.

Thus, In GCN devices, the four vector units within a CU can operate on four different wavefronts. If operations within a wavefront include dependencies, independent operations from different wavefronts can be selected to be assigned to a single vector unit to be executed in parallel every cycle.

2.2.2 Key differences between Southern Islands, Sea Islands, and Volcanic Islands families

The number of Asynchronous Compute Engines (ACEs) and CUs in an AMD GPU, and the way they are structured, vary with the device family, as well as with the device designations within a family.

The ACEs are responsible for managing the CUs and for scheduling and resource allocation of the compute tasks (but not of the graphics shader tasks). The ACEs operate independently; the greater the number of ACEs, the greater is the performance. Each ACE fetches commands from cache or memory, and creates task queues to be scheduled for execution on the CUs depending on their priority and on the availability of resources.

Each ACE contains up to 8 queues and, together with the graphics command processor, allows up to nine independent vector instructions to be executed per clock cycle. Some of these queues are not available for use by OpenCL.

Devices in the Southern Islands families typically have two ACEs. Devices in the Sea Islands and Volcanic Islands families contain between four and eight ACEs, so they offer more performance. For example, the AMD Radeon™ R9 290X devices, in the VI family contain 8 ACEs and 44 CUs.

2.3 Communication Between Host and the GPU Compute Device

The following subsections discuss the communication between the host (CPU) and the GPU in a compute device. This includes an overview of the PCIe bus, processing API calls, and DMA transfers.

Communication and data transfers between the system and the GPU compute device occur on the PCIe channel. AMD Accelerated Parallel Processing graphics cards use PCIe 2.0 x16 (second generation, 16 lanes). Generation 1 x16 has a theoretical maximum throughput of 4 GBps in each direction. Generation 2 x16 doubles the throughput to 8 GBps in each direction. Southern

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Islands AMD GPUs support PCIe 3.0 with a theoretical peak performance of 16 GBps. Actual transfer performance is CPU and chipset dependent.

Transfers from the system to the GPU compute device are done either by the *command processor* or by the *DMA engine*. The GPU compute device also can read and write system memory directly from the compute unit through kernel instructions over the PCIe bus.

2.3.1 Processing API Calls: The Command Processor

The host application does not interact with the GPU compute device directly. A driver layer translates and issues commands to the hardware on behalf of the application.

Most commands to the GPU compute device are buffered in a command queue on the host side. The queue of commands is sent to, and processed by, the GPU compute device. There is no guarantee as to when commands from the command queue are executed, only that they are executed in order.

Command queue elements include:

- Kernel execution calls
- Kernels
- Constants
- Transfers between device and host

2.3.2 DMA Transfers

Certain memory transfer calls use the DMA engine. To properly leverage the DMA engine, make the associated OpenCL data transfer calls. See the *AMD OpenCL Optimization Reference Guide* for more information.

Direct Memory Access (DMA) memory transfers can be executed separately from the command queue using the DMA engine on the GPU compute device. DMA calls are executed immediately; and the order of DMA calls and command queue flushes is guaranteed.

DMA transfers can occur asynchronously. This means that a DMA transfer is executed concurrently with other system or GPU compute operations when there are no dependencies. However, data is not guaranteed to be ready until the DMA engine signals that the event or transfer is completed. The application can use OpenCL to query the hardware for DMA event completion. If used carefully, DMA transfers are another source of parallelization.

Southern Island devices have two DMA engines that can perform bidirectional transfers over the PCIe bus with multiple queues created in consecutive order, since each DMA engine is assigned to an odd or an even queue correspondingly.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

2.3.3 Masking Visible Devices

By default, OpenCL applications are exposed to all GPUs installed in the system; this allows applications to use multiple GPUs to run the compute task.

In some cases, the user might want to mask the visibility of the GPUs seen by the OpenCL application. One example is to dedicate one GPU for regular graphics operations and the other three (in a four-GPU system) for Compute. To do that, set the `GPU_DEVICE_ORDINAL` environment parameter, which is a comma-separated list variable:

- Under Windows: set GPU_DEVICE_ORDINAL=1,2,3
- Under Linux: export GPU_DEVICE_ORDINAL=1,2,3

Another example is a system with eight GPUs, where two distinct OpenCL applications are running at the same time. The administrator might want to set `GPU_DEVICE_ORDINAL` to 0,1,2,3 for the first application, and 4,5,6,7 for the second application; thus, partitioning the available GPUs so that both applications can run at the same time.

2.4 Wavefront Scheduling

GPU compute devices are very efficient at parallelizing large numbers of work-items in a manner transparent to the application. Each GPU compute device uses the large number of wavefronts to hide memory access latencies by having the resource scheduler switch the active wavefront in a given compute unit whenever the current wavefront is waiting for a memory access to complete. Hiding memory access latencies requires that each work-item contain a large number of ALU operations per memory load/store.

Figure 2.6 shows the timing of a simplified execution of wavefronts in a single compute unit. At time 0, the wavefronts are queued and waiting for execution. In this example, only four wavefronts (T0…T3) are scheduled for the compute unit. The hardware limit for the number of active wavefront is dependent on the resource usage (such as the number of active registers used) of the program being executed. An optimally programmed GPU compute device typically has many of active wavefronts.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

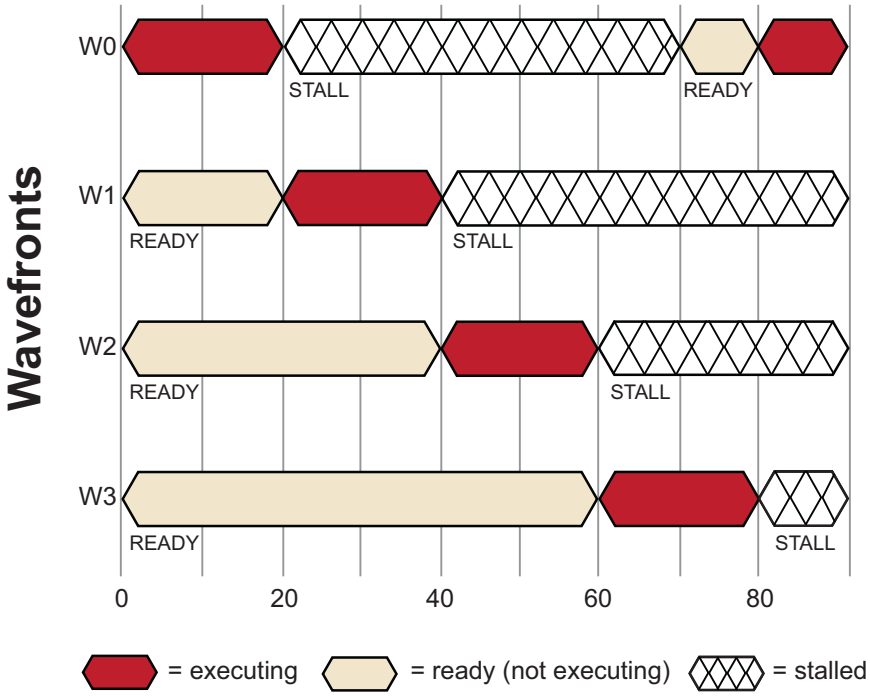

Figure 2.6 Simplified Execution Of Wavefront On A Single Compute Unit

At runtime, wavefront T0 executes until cycle 20; at this time, a stall occurs due to a memory fetch request. The scheduler then begins execution of the next wavefront, T1. Wavefront T1 executes until it stalls or completes. New wavefronts execute, and the process continues until the available number of active wavefronts is reached. The scheduler then returns to the first wavefront, T0.

If the data wavefront T0 is waiting for has returned from memory, T0 continues execution. In the example in Figure 2.6, the data is ready, so T0 continues. Since there were enough wavefronts and processing element operations to cover the long memory latencies, the compute unit does not idle. This method of memory latency hiding helps the GPU compute device achieve maximum performance.

If none of T0 – T3 are runnable, the compute unit waits (stalls) until one of T0 – T3 is ready to execute. In the example shown in Figure 2.7, T0 is the first to continue execution.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

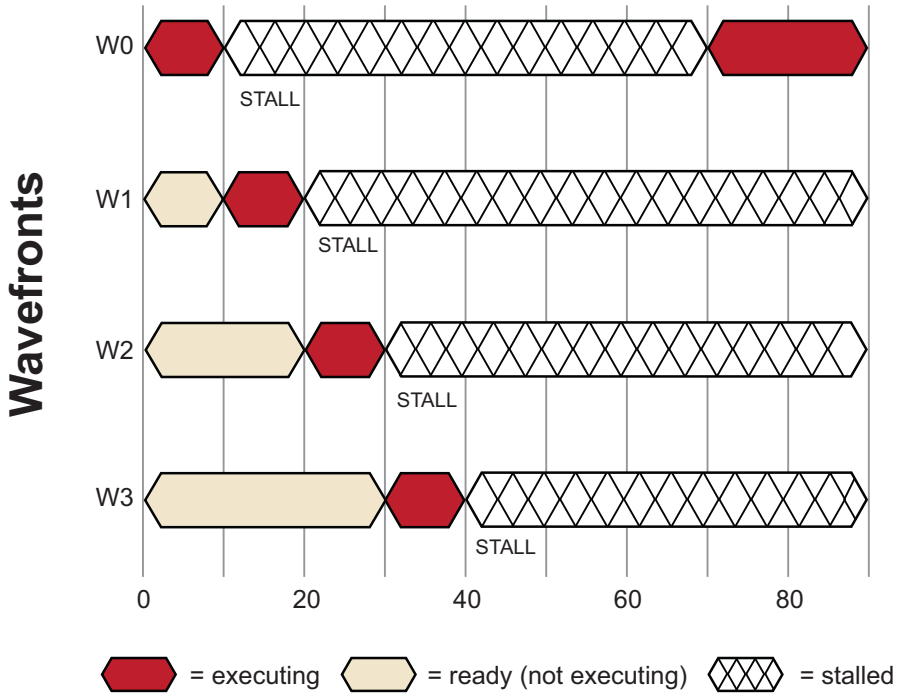

Figure 2.7 Compute Unit Stall Due to Data Dependency

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Chapter 3
Building and Running OpenCL Programs

The compiler tool-chain provides a common framework for both CPUs and GPUs, sharing the front-end and some high-level compiler transformations. The back-ends are optimized for the device type (CPU or GPU). Figure 3.1 is a high-level diagram showing the general compilation path of applications using OpenCL. Functions of an application that benefit from acceleration are re-written in OpenCL and become the OpenCL source. The code calling these functions are changed to use the OpenCL API. The rest of the application remains unchanged. The kernels are compiled by the OpenCL compiler to either CPU binaries or GPU binaries, depending on the target device.

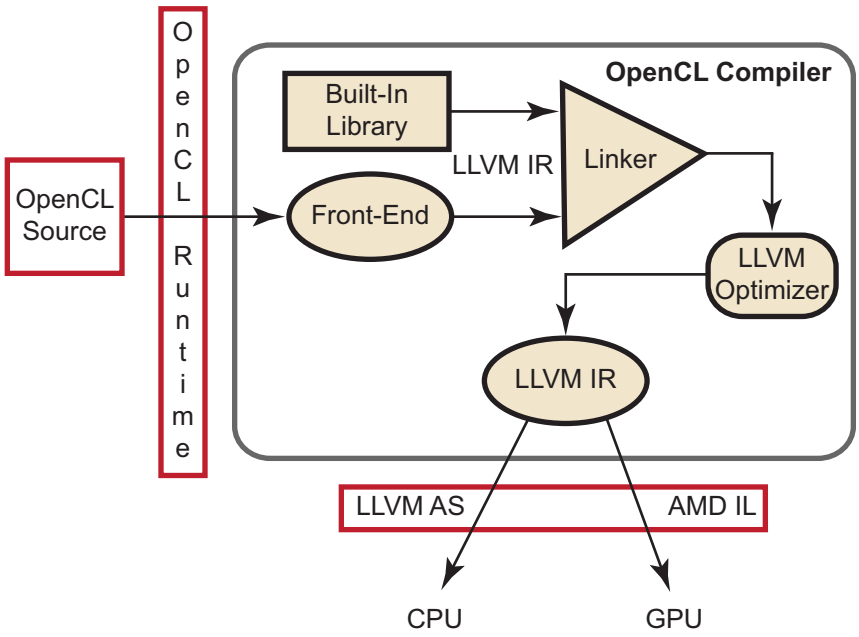

Figure 3.1 OpenCL Compiler Toolchain

For CPU processing, the OpenCL runtime uses the LLVM AS to generate x86 binaries. The OpenCL runtime automatically determines the number of processing elements, or cores, present in the CPU and distributes the OpenCL kernel between them.

For GPU processing, the OpenCL runtime post-processes the AMD IL from the OpenCL compiler and turns it into complete AMD IL. This adds macros (from a macro database, similar to the built-in library) specific to the GPU. The OpenCL

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Runtime layer then removes unneeded functions and passes the complete IL to the Shader compiler for compilation to GPU-specific binaries.

3.1 Compiling the Program

An OpenCL application consists of a host program (C/C++) and an optional kernel program (`.cl`). To compile an OpenCL application, the host program must be compiled; this can be done using an off-the-shelf compiler such as g++ or MSVC++. The application kernels are compiled into device-specific binaries using the OpenCL compiler.

This compiler uses a standard C front-end, as well as the low-level virtual machine (LLVM) framework, with extensions for OpenCL. The compiler starts with the OpenCL source that the user program passes through the OpenCL runtime interface (Figure 3.1). The front-end translates the OpenCL source to LLVM IR. It keeps OpenCL-specific information as metadata structures. (For example, to debug kernels, the front end creates metadata structures to hold the debug information; also, a pass is inserted to translate this into LLVM debug nodes, which includes the line numbers and source code mapping.) The front-end supports additional data-types (int4, float8, etc.), additional keywords (kernel, global, etc.) and built-in functions (`get_global_id()`, `barrier()`, etc.). Also, it performs additional syntactic and semantic checks to ensure the kernels meet the OpenCL specification. The input to the LLVM linker is the output of the front-end and the library of built-in functions. This links in the built-in OpenCL functions required by the source and transfers the data to the optimizer, which outputs optimized LLVM IR.

For GPU processing, the LLVM IR-to-AMD IL module receives LLVM IR and generates optimized IL for a specific GPU type in an incomplete format, which is passed to the OpenCL runtime, along with some metadata for the runtime layer to finish processing.

For CPU processing, LLVM AS generates x86 binary.

3.1.1 Compiling on Windows

To compile OpenCL applications on Windows requires that Visual Studio 2008 Professional Edition (or later) or the Intel C (C++) compiler are installed. All C++ files must be added to the project, which must have the following settings.

- Project Properties → C/C++ → Additional Include Directories
  These must include `$(AMDAPPSDKROOT)/include` for OpenCL headers. Optionally, they can include `$(AMDAPPSDKSAMPLESROOT)/include` for SDKUtil headers.
- Project Properties → C/C++ → Preprocessor Definitions
  These must define `AMD_OS_WIN`.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

- Project Properties → Linker → Additional Library Directories
  These must include `$(AMDAPPSDKROOT)/lib/x86` for OpenCL libraries. Optionally, they can include `$(AMDAPPSDKSAMPLESROOT)/lib/x86` for SDKUtil libraries.
- Project Properties → Linker → Input → Additional Dependencies
  These must include `OpenCL.lib`. Optionally, they can include `SDKUtil.lib`.

3.1.2 Compiling on Linux

To compile OpenCL applications on Linux requires that the gcc or the Intel C compiler is installed. There are two major steps to do this: compiling and linking.

1. Compile all the C++ files (`Template.cpp`), and get the object files.
   For 32-bit object files on a 32-bit system, or 64-bit object files on 64-bit system:

```
g++ -o Template.o -DAMD_OS_LINUX -c Template.cpp -I$AMDAPPSDKROOT/include
```

   For building 32-bit object files on a 64-bit system:

```
g++ -o Template.o -DAMD_OS_LINUX -c Template.cpp -I$AMDAPPSDKROOT/include
```

2. Link all the object files generated in the previous step to the OpenCL library and create an executable.

   For linking to a 64-bit library:

```
g++ -o Template Template.o -lOpenCL -L$AMDAPPSDKROOT/lib/x86_64
```

   For linking to a 32-bit library:

```
g++ -o Template Template.o -lOpenCL -L$AMDAPPSDKROOT/lib/x86
```

The OpenCL samples in the SDK provided by AMD Accelerated Parallel Processing depend on the SDKUtil library. In Linux, the samples use the shipped `SDKUtil.lib`, whether or not the sample is built for release or debug. When compiling all samples from the `samples/opencl` folder, the `SDKUtil.lib` is created first; then, the samples use this generated library. When compiling the SDKUtil library, the created library replaces the shipped library.

The following are linking options if the samples depend on the SDKUtil Library (assuming the SDKUtil library is created in `$AMDAPPSDKROOT/lib/x86_64` for 64-bit libraries, or `$AMDAPPSDKROOT/lib/x86` for 32-bit libraries).

```
g++ -o Template Template.o -lSDKUtil -lOpenCL -L$AMDAPPSDKROOT/lib/x86_64
```

```
g++ -o Template Template.o -lSDKUtil -lOpenCL -L$AMDAPPSDKROOT/lib/x86
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

3.1.3 Supported Standard OpenCL Compiler Options

The currently supported options are:

- `-I dir` — Add the directory `dir` to the list of directories to be searched for header files. When parsing `#include` directives, the OpenCL compiler resolves relative paths using the current working directory of the application.
- `-D name` — Predefine *name* as a macro, with definition = 1. For `-D name=definition`, the contents of *definition* are tokenized and processed as if they appeared during the translation phase three in a `#define` directive. In particular, the definition is truncated by embedded newline characters. `-D` options are processed in the order they are given in the *options* argument to `clBuildProgram`.

3.1.4 AMD-Developed Supplemental Compiler Options

The following supported options are not part of the OpenCL specification:

- `-g` — This is an experimental feature that lets you use the GNU project debugger, GDB, to debug kernels on x86 CPUs running Linux or cygwin/minGW under Windows. For more details, see Chapter 4, “Debugging OpenCL.” This option does not affect the default optimization of the OpenCL code.
- `-O0` — Specifies to the compiler not to optimize. This is equivalent to the OpenCL standard option `-cl-opt-disable`.
- `-f[no-]bin-source` — Does [not] generate OpenCL source in the `.source` section. For more information, see Appendix C, “OpenCL Binary Image Format (BIF) v2.0.”
- `-f[no-]bin-llvmir` — Does [not] generate LLVM IR in the `.llvmir` section. For more information, see Appendix C, “OpenCL Binary Image Format (BIF) v2.0.”
- `-f[no-]bin-amdil` — Does [not] generate AMD IL in the `.amdil` section. For more information, see Appendix C, “OpenCL Binary Image Format (BIF) v2.0.”
- `-f[no-]bin-exe` — Does [not] generate the executable (ISA) in `.text` section. For more information, see Appendix C, “OpenCL Binary Image Format (BIF) v2.0.”
- `-save-temps[=<prefix>]` — This option dumps intermediate temporary files, such as IL and ISA code, for each OpenCL kernel. If `<prefix>` is not given, temporary files are saved in the default temporary directory (the current directory for Linux, `C:\Users\<user>\AppData\Local` for Windows). If `<prefix>` is given, those temporary files are saved with the given `<prefix>`. If `<prefix>` is an absolute path prefix, such as `C:\your\work\dir\mydumpprefix`, those temporaries are saved under `C:\your\work\dir`, with `mydumpprefix` as prefix to all temporary names. For example,

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

`-save-temps`
under the default directory
`_temp_nn_xxx_yyy.il, _temp_nn_xxx_yyy.isa`

`-save-temps=aaa`
under the default directory
`aaa_nn_xxx_yyy.il, aaa_nn_xxx_yyy.isa`

`-save-temps=C:\you\dir\bbb`
under `C:\you\dir`
`bbb_nn_xxx_yyy.il, bbb_nn_xxx_yyy.isa`

where `xxx` and `yyy` are the device name and kernel name for this build, respectively, and `nn` is an internal number to identify a build to avoid overriding temporary files. Note that this naming convention is subject to change.

To avoid source changes, there are two environment variables that can be used to change CL options during the runtime.

- `AMD_OCL_BUILD_OPTIONS` — Overrides the CL options specified in `clBuildProgram()`.
- `AMD_OCL_BUILD_OPTIONS_APPEND` — Appends options to those specified in `clBuildProgram()`.

3.2 Running the Program

The runtime system assigns the work in the command queues to the underlying devices. Commands are placed into the queue using the `clEnqueue` commands shown in the listing below.

OpenCL API Function	Description
clCreateCommandQueue()	Create a command queue for a specific device (CPU, GPU).
clCreateProgramWithSource() clCreateProgramWithBinary()	Create a program object using the source code of the application kernels.
clBuildProgram()	Compile and link to create a program executable from the program source or binary.
clCreateKernel()	Creates a kernel object from the program object.
clCreateBuffer()	Creates a buffer object for use via OpenCL kernels.
clSetKernelArg() clEnqueueNDRangeKernel()	Set the kernel arguments, and enqueue the kernel in a command queue.
clEnqueueReadBuffer(), clEnqueueWriteBuffer()	Enqueue a command in a command queue to read from a buffer object to host memory, or write to the buffer object from host memory.
clEnqueueWaitForEvents()	Wait for the specified events to complete.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

The commands can be broadly classified into three categories.

- Kernel commands (for example, `clEnqueueNDRangeKernel()`, etc.),
- Memory commands (for example, `clEnqueueReadBuffer()`, etc.), and
- Event commands (for example, `clEnqueueWaitForEvents()`, etc.

As illustrated in Figure 3.2, the application can create multiple command queues (some in libraries, for different components of the application, etc.). These queues are muxed into one queue per device type. The figure shows command queues 1 and 3 merged into one CPU device queue (blue arrows); command queue 2 (and possibly others) are merged into the GPU device queue (red arrow). The device queue then schedules work onto the multiple compute resources present in the device. Here, K = kernel commands, M = memory commands, and E = event commands.

3.2.1 Running Code on Windows

The following steps ensure the execution of OpenCL applications on Windows.

1. The path to `OpenCL.lib` (`$AMDAPPSDKROOT/lib/x86`) must be included in path environment variable.
2. Generally, the path to the kernel file (`Template_Kernel.cl`) specified in the host program is relative to the executable. Unless an absolute path is specified, the kernel file must be in the same directory as the executable.

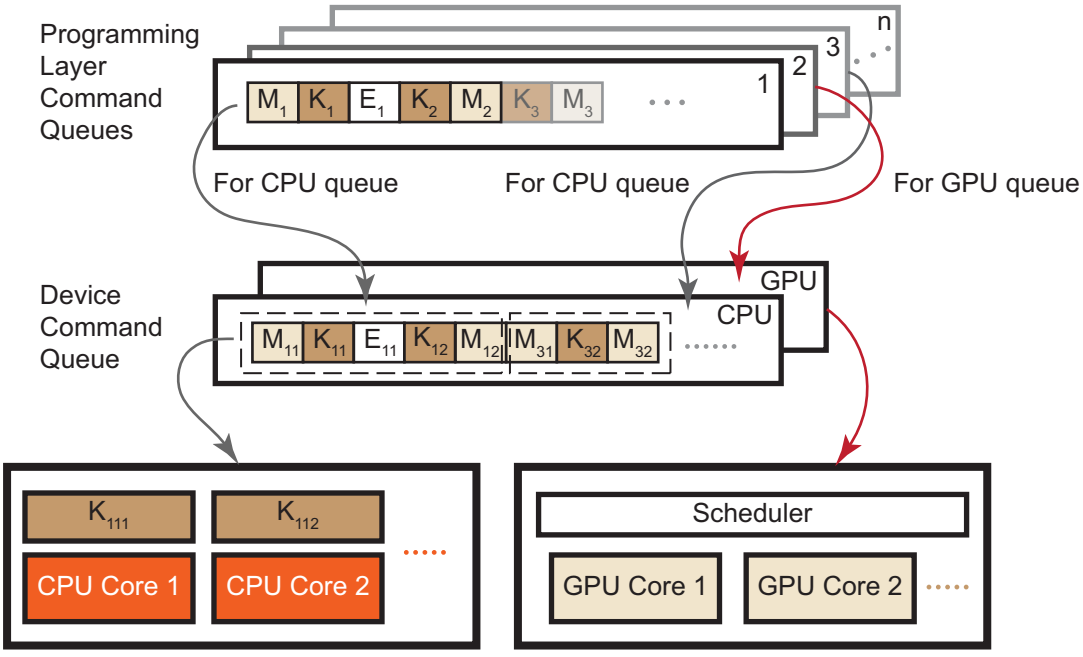

Figure 3.2 Runtime Processing Structure

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

3.2.2 Running Code on Linux

The following steps ensure the execution of OpenCL applications on Linux.

1. The path to `libOpenCL.so` (`$AMDAPPSDKROOT/lib/x86`) must be included in `$LD_LIBRARY_PATH`.
2. `/usr/lib/OpenCL/vendors/` must have `libatiocl32.so` and/or `libatiocl64.so`.
3. Generally, the path to the kernel file (`Template_Kernel.cl`) specified in the host program is relative to the executable. Unless an absolute path is specified, the kernel file must be in the same directory as the executable.

3.3 Calling Conventions

For all Windows platforms, the `__stdcall` calling convention is used. Function names are undecorated.

For Linux, the calling convention is `__cdecl`.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Chapter 4
Debugging OpenCL

This chapter discusses how to debug OpenCL programs running on AMD Accelerated Parallel Processing GPU and CPU compute devices. The preferred method is to debug with the AMD CodeXL, as described in Section 4.1, “AMD CodeXL GPU Debugger.” The second method, described in Section 4.2, “Debugging CPU Kernels with GDB,” is to use experimental features provided by AMD Accelerated Parallel Processing (GNU project debugger, GDB) to debug kernels on x86 CPUs running Linux or cygwin/minGW under Windows.

4.1 AMD CodeXL GPU Debugger

CodeXL 1.2 is available as an extension to Microsoft® Visual Studio®, a stand-alone version for Windows, and a stand-alone version for Linux. It offers real-time OpenCL kernel debugging and memory analysis on GPU devices and lets developers

- access the kernel execution directly from the API call that issues it,
- debug inside the kernel, and
- view all variable values across the different work-groups and work-items.

It also provides OpenGL API-level debugging and memory analysis. For information on downloading and installing CodeXL, see:

http://developer.amd.com/tools-and-sdks/heterogeneous-computing/codexl/

After installing CodeXL, launch Visual Studio, then open the solution to be worked on. In the Visual Studio menu bar, note the new CodeXL menu, which contains all the required controls.

Select a Visual C/C++ project, and set its debugging properties as normal. To add a breakpoint, either select New CodeXL Breakpoint from the CodeXL menu, or navigate to a kernel file (`*.cl`) used in the application, and set a breakpoint on the appropriate source line. Then, select the Launch OpenCL/OpenGL Debugging from the CodeXL menu to start debugging.

CodeXL currently supports only API-level debugging and OpenCL kernel debugging. Stepping through C/C++ code is not yet possible; however, the C/C++ call stack can be seen in the Visual Studio call stack view, which shows what led to the API function call.

To start kernel debugging, there are several options. One is to Step Into (F11) the appropriate `clEnqueueNDRangeKernel` function call. Once the kernel starts

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

executing, debug it like C/C++ code, stepping into, out of, or over function calls in the kernel, setting source breakpoints, and inspecting the locals, autos, watch, and call stack views.

If you develop on Linux, or do not use Visual Studio, using the CodeXL stand-alone application is very similar. After installing CodeXL, launch the CodeXL stand-alone application from the installation folder. On the start page, select "Create New Project," and use the browse button next to the "Executable Path" field to select your application. Click the "Go" (F5) button, and follow the instructions above to enter kernel debugging.

4.2 Debugging CPU Kernels with GDB

This section describes an experimental feature for using the GNU project debugger, GDB, to debug kernels on x86 CPUs running Linux or cygwin/minGW under Windows.

4.2.1 Setting the Environment

The OpenCL program to be debugged first is compiled by passing the "`-g -O0`" (or "`-g -cl-opt-disable`") option to the compiler through the options string to `clBuildProgram`. For example, using the C++ API:

```
err = program.build(devices,"-g -O0");
```

To avoid source changes, set the environment variable as follows:

```
AMD_OCL_BUILD_OPTIONS_APPEND="-g -O0" or
AMD_OCL_BUILD_OPTIONS="-g -O0"
```

Below is a sample debugging session of a program with a simple `hello world` kernel. The following GDB session shows how to debug this kernel. Ensure that the program is configured to be executed on the CPU. It is important to set `CPU_MAX_COMPUTE_UNITS=1`. This ensures that the program is executed deterministically.

4.2.2 Setting the Breakpoint in an OpenCL Kernel

To set a breakpoint, use:

```
b [N | function | kernel_name]
```

where *N* is the line number in the source code, *function* is the function name, and *kernel_name* is constructed as follows: if the name of the kernel is bitonicSort, the *kernel_name* is `__OpenCL_bitonicSort_kernel`.

Note that if no breakpoint is set, the program does not stop until execution is complete.

Also note that OpenCL kernel symbols are not visible in the debugger until the kernel is loaded. A simple way to check for known OpenCL symbols is to set a

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

breakpoint in the host code at `clEnqueueNDRangeKernel`, and to use the GDB `info functions __OpenCL` command, as shown in the example below.

4.2.3 Sample GDB Session

The following is a sample debugging session. Note that two separate breakpoints are set. The first is set in the host code, at `clEnqueueNDRangeKernel()`. The second breakpoint is set at the actual CL kernel function.

```
$ export AMD_OCL_BUILD_OPTIONS_APPEND="-g -O0"
$ export CPU_MAX_COMPUTE_UNITS=1
$ gdb BitonicSort
GNU gdb (GDB) 7.1-ubuntu
Copyright (C) 2010 Free Software Foundation, Inc.
License GPLv3+: GNU GPL version 3 or later
<http://gnu.org/licenses/gpl.html>
This is free software: you are free to change and redistribute it.
There is NO WARRANTY, to the extent permitted by law.  Type "show copying"
and "show warranty" for details.
This GDB was configured as "x86_64-linux-gnu".
For bug reporting instructions, please see:
<http://www.gnu.org/software/gdb/bugs/>...
Reading symbols from /home/himanshu/Desktop/ati-stream-sdk-v2.3-
lnx64/samples/opencl/bin/x86_64/BitonicSort...done.
(gdb) b clEnqueueNDRangeKernel
Breakpoint 1 at 0x403228
(gdb) r --device cpu
Starting program: /home/himanshu/Desktop/ati-stream-sdk-v2.3-
lnx64/samples/opencl/bin/x86_64/BitonicSort --device cpu
[Thread debugging using libthread_db enabled]

Unsorted Input
53 5 199 15 120 9 71 107 71 242 84 150 134 180 26 128 196 9 98 4 102 65
206 35 224 2 52 172 160 94 2 214 99 .....


Platform Vendor : Advanced Micro Devices, Inc.
Device 0 : AMD Athlon(tm) II X4 630 Processor
[New Thread 0x7ffff7e6b700 (LWP 1894)]
[New Thread 0x7ffff2fcc700 (LWP 1895)]
Executing kernel for 1 iterations
-----------------------------------------

Breakpoint 1, 0x00007ffff77b9b20 in clEnqueueNDRangeKernel () from
/home/himanshu/Desktop/ati-stream-sdk-v2.3-lnx64/lib/x86_64/libOpenCL.so
(gdb) info functions __OpenCL
All functions matching regular expression "__OpenCL":

File OCLm2oVFr.cl:
void __OpenCL_bitonicSort_kernel(uint *, const uint, const uint, const
uint, const uint);

Non-debugging symbols:
0x00007ffff23c2dc0  __OpenCL_bitonicSort_kernel@plt
0x00007ffff23c2f40  __OpenCL_bitonicSort_stub
(gdb) b __OpenCL_bitonicSort_kernel
Breakpoint 2 at 0x7ffff23c2de9: file OCLm2oVFr.cl, line 32.
(gdb) c
Continuing.
[Switching to Thread 0x7ffff2fcc700 (LWP 1895)]

Breakpoint 2, __OpenCL_bitonicSort_kernel (theArray=0x615ba0, stage=0,
passOfStage=0, width=1024, direction=0) at OCLm2oVFr.cl:32
32      uint sortIncreasing = direction;
(gdb) p get_global_id(0)
$1 = 0
(gdb) c
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
Continuing.

Breakpoint 2, __OpenCL_bitonicSort_kernel (theArray=0x615ba0, stage=0,
passOfStage=0, width=1024, direction=0) at OCLm2oVFr.cl:32
32      uint sortIncreasing = direction;
(gdb) p get_global_id(0)
$2 = 1
(gdb)
```

4.2.4 Notes

1. To make a breakpoint in a working thread with some particular ID in dimension N, one technique is to set a conditional breakpoint when the `get_global_id(N) == ID`. To do this, use:

   ```
   b [ N | function | kernel_name ] if (get_global_id(N)==ID)
   ```

   where `N` can be 0, 1, or 2.

2. For complete GDB documentation, see http://www.gnu.org/software/gdb/documentation/ .
3. For debugging OpenCL kernels in Windows, a developer can use GDB running in cygwin or minGW. It is done in the same way as described in sections 3.1 and 3.2.

   Notes:

   – Only OpenCL kernels are visible to GDB when running cygwin or minGW. GDB under cygwin/minGW currently does not support host code debugging.

   – It is not possible to use two debuggers attached to the same process. Do not try to attach Visual Studio to a process, and concurrently GDB to the kernels of that process.

   – Continue to develop the application code using Visual Studio. Currently, gcc running in cygwin or minGW is not supported.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Chapter 5
OpenCL Static C++ Programming Language

5.1 Overview

This extension defines the OpenCL Static C++ kernel language, which is a form of the ISO/IEC Programming languages C++ specification[1]. This language supports overloading and templates that can be resolved at compile time (hence static), while restricting the use of language features that require dynamic/runtime resolving. The language also is extended to support most of the features described in Section 6 of OpenCL spec: new data types (vectors, images, samples, etc.), OpenCL Built-in functions, and more.

5.1.1 Supported Features

The following list contains the major C++ features supported by this extension.

- Kernel and function overloading.
- Inheritance:
  - Strict inheritance.
  - Friend classes.
  - Multiple inheritance.
- Templates:
  - Kernel templates.
  - Member templates.
  - Template default argument.
  - Limited class templates (the `virtual`. keyword is not exposed).
  - Partial template specialization
- Namespaces.
- References.
- `this` operator.

Note that supporting templates and overloading highly improve the efficiency of writing code: it allows developers to avoid replication of code when not necessary.

1. Programming languages C++. International Standard ISO/IEC 14881, 1998.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Using kernel template and kernel overloading requires support from the runtime API as well. AMD provides a simple extension to `clCreateKernel`, which enables the user to specify the desired kernel.

5.1.2 Unsupported Features

C++ features not supported by this extension are:

- Virtual functions (methods marked with the `virtual` keyword).
- Abstract classes (a class defined only of pure virtual functions).
- Dynamic memory allocation (non-placement new/delete support is not provided).
- Exceptions (no support for throw/catch).
- The :: operator.
- STL and other standard C++ libraries.
- The language specified in this extension can be easily expanded to support these features.

5.1.3 Relations with ISO/IEC C++

This extension focuses on documenting the differences between the OpenCL Static C++ kernel language and the ISO/IEC Programming languages C++ specification. Where possible, this extension leaves technical definitions to the ISO/IEC specification.

5.2 Additions and Changes to Section 5 - The OpenCL C Runtime

5.2.1 Additions and Changes to Section 5.7.1 - Creating Kernel Objects

In the static C++ kernel language, a kernel can be overloaded, templated, or both. The syntax explaining how to do it is defined in Sections 5.3.4 and 5.3.5, below.

To support these cases, the following error codes were added; these can be returned by `clCreateKernel`.

- `CL_INVALID_KERNEL_TEMPLATE_TYPE_ARGUMENT_AMD` if a kernel template argument is not a valid type (is neither a valid OpenCL C type or a user defined type in the same source file).
- `CL_INVALID_KERNEL_TYPE_ARGUMENT_AMD` if a kernel type argument, used for overloading resolution, is not a valid type (is neither a valid OpenCL C type or user-defined type in the same source program).

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

5.2.2 Passing Classes between Host and Device

This extension allows a developer to pass classes between the host and the device. The mechanism used to pass the class to the device and back are the existing buffer object APIs. The class that is passed maintains its state (public and private members), and the compiler implicitly changes the class to use either the host-side or device-side methods.

On the host side, the application creates the class and an equivalent memory object with the same size (using the `sizeof` function). It then can use the class methods to set or change values of the class members. When the class is ready, the application uses a standard buffer API to move the class to the device (either Unmap or Write), then sets the buffer object as the appropriate kernel argument and enqueues the kernel for execution. When the kernel finishes the execution, the application can map back (or read) the buffer object into the class and continue working on it.

5.3 Additions and Changes to Section 6 - The OpenCL C Programming Language

5.3.1 Building C++ Kernels

To compile a program that contains C++ kernels and functions, the application must add the following compile option to `clBuildProgramWithSource`:

```
-x language
```

where `language` is defined as one of the following:

- `clc` – the source language is considered to be OpenCL C, as defined in the The OpenCL Programming Language version 1.2[1].
- `clc++` - the source language is considered to be OpenCL C++, as defined in the following sections of the this document.

5.3.2 Classes and Derived Classes

OpenCL C is extended to support classes and derived classes as per Sections 9 and 10 of the C++ language specification, with the limitation that virtual functions and abstracts classes are not supported. The virtual keyword is reserved, and the OpenCL C++ compiler is required to report a compile time error if it is used in the input program.

This limitation restricts class definitions to be fully statically defined. There is nothing prohibiting a future version of OpenCL C++ from relaxing this restriction, pending performance implications.

A class definition can not contain any address space qualifier, either for members or for methods:

1. *The OpenCL Programming Language 1.2. Rev15*, Khronos 2011.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
class myClass{
    public:
        int myMethod1(){ return x;}
        void __local myMethod2(){x = 0;}
    private:
        int x;
        __local y; // illegal
};
```

The class invocation inside a kernel, however, can be either in private or local address space:

```
__kernel void myKernel()
{
    myClass c1;
    __local myClass c2;
    ...
}
```

Classes can be passed as arguments to kernels, by defining a buffer object at the size of the class, and using it. The device invokes the adequate device-specific methods, and accesses the class members passed from the host.

OpenCL C kernels (defined with `__kernel`) may not be applied to a class constructor, destructor, or method, except in the case that the class method is defined static and thus does not require object construction to be invoked.

5.3.3 Namespaces

Namespaces are support without change as per [1].

5.3.4 Overloading

As defined in of the C++ language specification, when two or more different declarations are specified for a single name in the same scope, that name is said to be overloaded. By extension, two declarations in the same scope that declare the same name but with different types are called overloaded declarations. Only kernel and function declarations can be overloaded, not object and type declarations.

As per of the C++ language specification, there are a number of restrictions as to how functions can be overloaded; these are defined formally in Section 13 of the C++ language specification. Note that kernels and functions cannot be overloaded by return type.

Also, the rules for well-formed programs as defined by Section 13 of the C++ language specification are lifted to apply to both kernel and function declarations.

The overloading resolution is per Section 13.1 of the C++ language specification, but extended to account for vector types. The algorithm for “best viable function”, Section 13.3.3 of the C++ language specification, is extended for vector types by inducing a partial-ordering as a function of the partial-ordering of its elements. Following the existing rules for vector types in the OpenCL 1.2 specification, explicit conversion between vectors is not allowed. (This reduces the number of

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

possible overloaded functions with respect to vectors, but this is not expected to be a particular burden to developers because explicit conversion can always be applied at the point of function evocation.)

For overloaded kernels, the following syntax is used as part of the kernel name:

foo(type$_1$,...,type$_n$)

where type$_1$,...,type$_n$ must be either an OpenCL scalar or vector type, or can be a user-defined type that is allocated in the same source file as the kernel foo.

To allow overloaded kernels, use the following syntax:

```
__attribute__((mangled_name(myMangledName)))
```

The kernel mangled_name is used as a parameter to pass to the clCreateKernel() API. This mechanism is needed to allow overloaded kernels without changing the existing OpenCL kernel creation API.

5.3.5 Templates

OpenCL C++ provides unrestricted support for C++ templates, as defined in Section 14 of the C++ language specification. The arguments to templates are extended to allow for all OpenCL base types, including vectors and pointers qualified with OpenCL C address spaces (i.e. __global, __local, __private, and __constant).

OpenCL C++ kernels (defined with __kernel) can be templated and can be called from within an OpenCL C (C++) program or as an external entry point (from the host).

For kernel templates, the following syntax is used as part of the kernel name (assuming a kernel called foo):

foo<type$_1$,...,type$_n$>

where type$_1$,...,type$_n$ must be either OpenCL scalar or vector type, or can be a user-defined type that is allocated in the same source file as the kernel foo. In this case a kernel is both overloaded and templated:

foo<type$_1$,...,type$_n$>(type$_{n+1}$,...,type$_m$)

Note that here overloading resolution is done by first matching non-templated arguments in order of appearance in the definition, then substituting template parameters. This allows intermixing of template and non-template arguments in the signature.

To support template kernels, the same mechanism for kernel overloading is used. Use the following syntax:

```
__attribute__((mangled_name(myMangledName)))
```

The kernel mangled_name is used as a parameter to passed to t he clCreateKernel() API. This mechanism is needed to allow template kernels

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

without changing the existing OpenCL kernel creation API. An implementation is not required to detect name collision with the user-specified `kernel_mangled` names involved.

5.3.6 Exceptions

Exceptions, as per Section 15 of the C++ language specification, are not supported. The keywords `try`, `catch`, and `throw` are reserved, and the OpenCL C++ compiler must produce a static compile time error if they are used in the input program.

5.3.7 Libraries

Support for the general utilities library, as defined in Sections 20-21 of the C++ language specification, is not provided. The standard C++ libraries and STL library are not supported.

5.3.8 Dynamic Operation

Features related to dynamic operation are not supported:

- the `virtual` modifier.
  OpenCL C++ prohibits the use of the `virtual` modifier. Thus, virtual member functions and virtual inheritance are not supported.
- Dynamic cast that requires runtime check.
- Dynamic storage allocation and deallocation.

5.4 Examples

5.4.1 Passing a Class from the Host to the Device and Back

The class definition must be the same on the host code and the device code, besides the members' type in the case of vectors. If the class includes vector data types, the definition must conform to the table that appears on Section 6.1.2 of the OpenCL Programming Specification 1.2, Corresponding API type for OpenCL Language types.

Example Kernel Code

```
Class Test
{
   setX (int value);
   private:
   int x;
}

__kernel foo (__global Test* InClass, ...)
{
   If (get_global_id(0) == 0)
      InClass->setX(5);
}
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Example Host Code

```
Class Test
{
   setX (int value);
   private:
   int x;
}


MyFunc ()
{
   tempClass = new(Test);
   ... // Some OpenCL startup code – create context, queue, etc.
                       cl_mem classObj = clCreateBuffer(context,
                       CL_MEM_USE_HOST_PTR, sizeof(Test),
                       &tempClass, event);
   clEnqueueMapBuffer(...,classObj,...);
   tempClass.setX(10);
   clEnqueueUnmapBuffer(...,classObj,...); //class is passed to the Device
   clEnqueueNDRange(..., fooKernel, ...);
   clEnqueueMapBuffer(...,classObj,...); //class is passed back to the Host

}
```

5.4.2 Kernel Overloading

This example shows how to define and use `mangled_name` for kernel overloading, and how to choose the right kernel from the host code. Assume the following kernels are defined:

```
__attribute__((mangled_name(testAddFloat4))) kernel void
testAdd(global float4 * src1, global float4 * src2, global float4 * dst)
{
   int tid = get_global_id(0);
   dst[tid] = src1[tid] + src2[tid];

}

__attribute__((mangled_name(testAddInt8))) kernel void
   testAdd(global int8 * src1, global int8 * src2, global int8 * dst)
{
   int tid = get_global_id(0);
   dst[tid] = src1[tid] + src2[tid];

}
```

The names `testAddFloat4` and `testAddInt8` are the external names for the two kernel instants. When calling `clCreateKernel`, passing one of these kernel names leads to the correct overloaded kernel.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

5.4.3 Kernel Template

This example defines a kernel template, `testAdd`. It also defines two explicit instants of the kernel template, `testAddFloat4` and `testAddInt8`. The names `testAddFloat4` and `testAddInt8` are the external names for the two kernel template instants that must be used as parameters when calling to the `clCreateKernel` API.

```
template <class T>
   kernel void testAdd(global T * src1, global T * src2, global T * dst)
{
   int tid = get_global_id(0);
   dst[tid] = src1[tid] + src2[tid];
}

template __attribute__((mangled_name(testAddFloat4))) kernel void
   testAdd(global float4 * src1, global float4 * src2, global float4 *
dst);

template __attribute__((mangled_name(testAddInt8))) kernel void
   testAdd(global int8 * src1, global int8 * src2, global int8 * dst);
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Chapter 6
OpenCL 2.0

6.1 Introduction

The OpenCL 2.0 specification is a significant evolution of OpenCL. It introduces features that allow closer collaboration between the host and OpenCL devices, such as Shared Virtual Memory (SVM) and device-side enqueue. Other features, such as pipes, dynamic parallelism, and new image-related additions provide effective ways of expressing heterogeneous programming constructs.

The following sections highlight the salient features of OpenCL 2.0 and provide usage guidelines.

- Shared Virtual Memory (SVM)
- Generic Address Space
- Device-side enqueue
- Atomics and synchronization
- Pipes
- Sub-groups
- Program-scope global Variables
- Image Enhancements
- Non-uniform work group size

Sample code is included wherever appropriate; complete samples illustrating the OpenCL 2.0 features are provided with the AMD APP SDK.

For guidelines on how to migrate from OpenCL 1.2 to OpenCL 2.0 and for information about querying for image- and device-specific extensions, see Portability considerations.

For a list of the new and deprecated functions, see Appendix F, “New and deprecated functions in OpenCL 2.0.”

6.2 Shared Virtual Memory (SVM)

6.2.1 Overview

In OpenCL 1.2, the host and OpenCL devices do not share the same virtual address space. Consequently, the host memory, the device memory, and communication between the host and the OpenCL devices, need to be explicitly

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

specified and managed. Buffers may need to be copied over to the OpenCL device memory for processing and copied back after processing; the same memory cannot be accessed simultaneously by the host and the OpenCL devices. To access locations within a buffer (or regions within an image), the appropriate offsets must be passed to and from the OpenCL devices.

In OpenCL 2.0, the host and OpenCL devices may share the same virtual address space. Buffers need not be copied over between devices. When the host and the OpenCL devices share the address space, communication between the host and the devices can occur via shared memory. This simplifies programming in heterogeneous contexts.

Support for SVM does not imply or require that the host and the OpenCL devices in an OpenCL 2.0 compliant architecture share actual physical memory. The OpenCL runtime manages the transfer of data between the host and the OpenCL devices; the process is transparent to the programmer, who sees a unified address space.

A caveat, however, concerns situations in which the host and the OpenCL devices access the same region of memory at the same time. It would be highly inefficient for the host and the OpenCL devices to have a consistent view of the memory for each load/store from any device/host. In general, the memory model of the language or architecture implementation determines how or when a memory location written by one thread or agent is visible to another. The memory model also determines to what extent the programmer can control the scope of such accesses.

OpenCL 2.0 adopts the memory model defined C++11 with some extensions. The memory orders taken from C++11 are: "relaxed", "acquire", "release", “acquire-release”, and "sequential consistent".

OpenCL 2.0 introduces a new (C++11-based) set of atomic operations with specific memory-model based semantics. Atomic operations are indivisible: a thread or agent cannot see partial results. The atomic operations supported are:

- `atomic_load/store`
- `atomic_init`
- `atomic_work_item_fence`
- `atomic_exchange`
- `atomic_compare_exchange`
- `atomic_fetch_<op>`, where `<op>` is "`add`", "`sub`", "`xor`", "`and`", or "`or`"

OpenCL 2.0 introduces the concept of "memory scope", which limits the extent to which atomic operations are visible. For example:

- "workgroup" scope means that the updates are to be visible only within the work group
- "device" scope means that the updates are to be visible only within the device (across workgroups within the device)

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

- "all-svm-devices" scope means the updates are available across devices (GPUs and the host/CPU).

The OpenCL 2.0 further differentiates between coarse-grained and fine-grained (buffer and system) SVM buffer sharing mechanisms. These mechanisms define the granularity at which the SVM buffers are shared.

Updates to coarse-grained or fine-grained SVM are visible to other devices at synchronization points:

- For coarse-grained SVM, the synchronization points are: the mapping or un-mapping of the SVM memory and kernel launch or completion. This means that any updates are visible only at the end of the kernel or at the point of un-mapping the region of memory. Coarse-grained buffer memory has a fixed virtual address for all devices it is allocated on. The physical memory is allocated on Device Memory.
- For fine-grained SVM, the synchronization points include, in addition to those defined for coarse-grained SVM, the mapping or un-mapping of memory and atomic operations. This means that updates are visible at the level of atomic operations on the SVM buffer (for fine-grained buffer SVM, allocated with the `CL_MEM_SVM_ATOMICS` flag) or the SVM system, i.e. anywhere in the SVM (for fine-grained system SVM). Fine-grained buffer memory has the same virtual address for all devices it is allocated on. The physical memory is allocated on the Device-Visible Host Memory. If the fine grain buffer is allocated with the `CL_MEM_SVM_ATOMICS` flag, the memory will be GPU-CPU coherent.

The OpenCL 2.0 specification mandates coarse-grained SVM but not fine-grained SVM.

For details, the developer is urged to read Section 3.3 of the OpenCL 2.0 specification.

6.2.2 Usage

In OpenCL 2.0, SVM buffers shared between the host and OpenCL devices are created by calling `clSVMAlloc` (or `malloc`/`new` in the case of fine-grain system support). The contents of such buffers may include pointers (into SVM buffers). Pointer-based data structures are especially useful in heterogenous programming scenarios. A typical scenario is as follows:

1. Host creates SVM buffer(s) with `clSVMAlloc`
2. Host maps the SVM buffer(s) with blocking `clEnqueueSVMMap`
3. Host fills/updates the SVM buffer(s) with data structures including pointers
4. Host unmaps the SVM buffer(s) by using `clEnqueueSVMUnmap`
5. Host enqueues processing kernels, passing SVM buffers to the kernels with calls to `clSetKernelArgSVMPointer` and/or `clSetKernelExecInfo`

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

6. The OpenCL 2.0 device processes the structures in SVM buffer(s) including following/updating pointers.

7. Repeat Step 2 as necessary.

Note that the map and unmap operations in Steps 2 and 4 may be eliminated if the SVM buffers are created by using the `CL_MEM_SVM_FINE_GRAIN_BUFFER` flag, which may not be supported on all devices.

6.2.2.1 Coarse-grained memory

Some applications do not require fine-grained atomics to ensure that the SVM is consistent across devices after each read/write access. After the initial map/creation of the buffer, the GPU or any other devices typically read from memory. Even if the GPU or other devices write to memory, they may not require a consistent view of the memory.

For example, while searching in parallel on a binary search tree (as in the Binary Search tree sample presented later in this section), coarse-grain buffers are usually sufficient. In general, coarse-grain buffers provide faster access compared to fine grain buffers as the memory is not required to be consistent across devices.

```
    for (i = 0; i < keys_per_wi; i++) {
           key = search_keys[init_id + i];
           tmp_node = root;
           while (1) {
                  if (!tmp_node || (tmp_node->value == key))
                         break;
                  tmp_node = (key < tmp_node->value) ? tmp_node-
>left : tmp_node->right;
           }
           found_nodes[init_id + i] = tmp_node;
    }
```

In the above example, the binary search tree root is created using coarse-grain SVM on the host as

```
           int flags = CL_MEM_READ_ONLY;
           if ((data = (node *) clSVMAlloc(context, flags, num_nodes
* sizeof(node), 0)) == NULL) {
                  printf("Error allocating memory for nodes.\n");
                  exit(1);
           }
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

The “data” is the tree created by the host as a coarse-grain buffer and is passed to the kernel as an input pointer.

Tree (size in M)	CPU – time (ms)	GPU (OpenCL 2.0)	GPU (OpenCL 1.2) *
1	23.46	5.17	3.22 (+1.92) (+49.58 +8.50)
5	86.11	24.87	13.95 (+4.15) (+259.12 +34.30)
10	180.73	51.58	N/A
25	381.77	129.58	N/A

The above table shows the performance of the 2.0 implementation over the 1.2 implementation.

As SVM was absent in OpenCL 1.2, the tree needed to be sent to the GPU memory. In addition, as pointers were not pointing to the same tree as on the host, the tree might have required to be transformed into arrays for GPU compute. The values provided in parentheses indicate the extra time required for such computes. Alternatively, offsets might have been required to be used instead of pointers on both the host and the device at the cost of a few more additions.

Finally, more than 5M nodes could not be allocated in 1.2, as the allowable memory allocation was limited by the amount of memory that could be used on the device. Overall, the 2.0 version exceeds the 1.2 version in both performance and usability.

6.3 Generic Address Space

6.3.1 Overview

In OpenCL 1.2, all parameters in a function definition must have address spaces associated with them. (The default address space is the private address space.) This necessitates creating an explicit version of the function must be created for each desired address space.

OpenCL 2.0 introduces a new address space called the generic address space. Data cannot be stored in the generic address space, but a pointer to this space can reference data located in the private, local, or global address spaces. A function with generic pointer arguments may be called with pointers to any address space except the constant address space. Pointers that are declared without pointing to a named address space point to the generic address space.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

However, such pointers must be associated with a named address space before they can be used. Functions may be written with arguments and return values that point to the generic address space, improving readability and programmability.

6.3.2 Usage

6.3.2.1 Generic example

In OpenCL 1.2, the developer needed to write three functions for a pointer `p` that can reference the local, private, or global address space:

```
void fooL (local int *p) { … }

void fooP (private int *p) { … }

void fooG (global int *p) { … }
```

In OpenCL 2.0, the developer needs to write only one function:

```
void foo (int *p)
```

As `foo` is a generic function, the compiler will accept calls to it with pointers to any address space except the constant address space.

Note: The OpenCL 2.0 spec itself shows most built-in functions that accept pointer arguments as accepting generic pointer arguments.

6.3.2.2 AMD APP SDK example

In the *xxx* APP SDK sample, `addMul2d` is a generic function that uses generic address spaces for its operands. The function computes the convolution sum of two vectors. Two kernels compute the convolution: one uses data in the `global` address space (`convolution2DUsingGlobal`); the other uses the `local` address space (`sepiaToning2DUsingLocal`). The use of a single function improves the readability of the source.

```
float4 addMul2D (uchar4 *src, float *filter, int2 filterDim, int
width)

{   int i, j;

    float4 sum = (float4)(0);

    for(i = 0; i < (filterDim.y); i++)

    {

           for(j = 0; j < (filterDim.x); j++)

           {

                  sum +=
(convert_float4(src[(i*width)+j]))*((float4)(filter[(i*filterDim.x)
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
+j]));
            }
    }
    return sum;
}
```

Note: The compiler will try to resolve the address space at compile time. Otherwise, the runtime will decide whether the pointer references the local or the global address space. For optimum performance, the code must make it easy for the compiler to detect the pointer reference by avoiding data-dependent address space selection, so that run-time resolution -- which is costly -- is not required.

6.4 Device-side enqueue

6.4.1 Overview

In OpenCL 1.2, a kernel cannot be enqueued from a currently running kernel. Enqueuing a kernel requires returning control to the host.

OpenCL 2.0 introduces "clang blocks" and new built-in functions that allow kernels to enqueue new work to device-side queues. In addition, the runtime API call `clCreateCommandQueue` has been deprecated in favor of a new call, `clCreateCommandQueueWithProperties`, which allows creating device-side command queues. Device enqueue allows programmers to launch the (child) kernels from the current (parent) kernel.

This mechanism offers the following benefits:

1. As the overhead of returning control to the host to launch the kernel is eliminated, device enqueue can improve the performance of applications. Some platforms (such as AMD's) provide a standard way of enqueuing work to the hardware, which can further improve the performance. Device-side enqueue has been observed to reduce by the overhead of enqueuing by more than 3x in some cases. Device-side enqueue can especially benefit applications that are inherently recursive or comprise work to be performed dynamically.
2. In OpenCL 1.2, knowing when all the workgroups of the running kernel have completed execution requires waiting on a completion event from that kernel. This necessitates waiting on the host before another kernel can be launched, if the host requires the result of a computation before it can proceed. Device-side enqueue allows the parent kernel to launch the child kernel and thus eliminates this wait. The *extracting primes* example in the subsequent sections illustrate this benefit.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

6.4.2 Usage

6.4.2.1 Iterate until convergence

Suppose a complex process requires 4 kernels, A, B, C, and Check, and that these kernels must be run in order repeatedly until the Check kernel produces a value indicating that the process has converged.

In OpenCL 1.2, the host side code to perform this might be structured as follows:

1. Enqueue kernel A
2. Enqueue kernel B
3. Enqueue kernel C
4. Enqueue kernel Check
5. Enqueue blocking map of Check result, e.g. with `clEnqueueSVMMap`
6. If Check result is not "Converged" then:
   Enqueue unmap of Check result
7. Go to Step 1

However, with device-side enqueue in OpenCL 2.0, the Check kernel may be altered to enqueue blocks that carry out A, B, C, and Check when it detects that convergence has not been reached. This avoids a potentially costly interaction with the host on each iteration. Also, a slight modification of Check might allow the replacement of the entire loop above with a single host-side enqueue of the Check kernel.

6.4.2.2 Data-dependent refinement

Consider a search or computational process that works from coarse levels to increasingly finer levels that operates something like this:

1. Search/Compute over current region
2. Loop over sub-regions in current region
3. If a sub-region is interesting:
    1. Refine the sub-region
    2. Apply a process to the refined sub-region

With OpenCL 1.2, this process would require a complex interaction between the host and the OpenCL device. The device-side kernel would need to somehow mark the sub-regions requiring further work, and the host side code would need to scan all of the sub-regions looking for the marked ones and then enqueue a kernel for each marked sub-region. This is made more difficult by the lack of globally visible atomic operations in OpenCL 1.2.

However, with OpenCL 2.0, rather than just marking each interesting sub-region, the kernel can instead launch a new sub-kernel to process each marked sub-

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

region. This significantly simplifies the code and improves efficiency due to the elimination of the interactions with, and dependence on, the host.

6.4.2.3 Extracting Primes from an array by using device-side enqueue

Given an array of positive integers, the device-side enqueue feature in OpenCL 2.0 can be used to efficiently extract the prime numbers present in that array and store them into another output array.

The following figure explains the input and output expected from this sample:

Input:

12	31	47	64	19	27	49	81	99	11

Output:

31	47	19	11

As can be seen in the Figure above, the input is given as an array of positive numbers and the expected output is an array of the primes in the input array. This is a classic example of data-parallel processing, in which the processing for each array element for checking whether the element is prime or not, can be done in parallel with all others.

In OpenCL 1.2, this processing can be done as follows:

```
kernel extractPrimes(global int *in, global int *out) {

int id = get_global_id (0);

if (isPrime(in[id]))

          out[id] = in[id];

}
```

Assuming that the output array is initialized with zeroes before the kernel executes, the output would be:

0	31	47	0	19	0	0	0	0	11

The output array can be written into by using atomics directly. But doing so would essentially make the loop run sequentially.

Extracting the non-zero numbers will require atomics or processing on the CPU through a sequential loop, both of which may not be inefficient or scalable.

In OpenCL 2.0, the new workgroup built-ins can efficiently compute prefix sums and perform reductions on various operations like additions, maximum and minimum on a workgroup. Some such built-in are:
`work_group_scan_inclusive_<op>`,
`work_group_scan_exclusive_<op>`, and `work_group_reduce_<op>`, where `<op>` can be one of "`add`", "`max`", or "`min`". The "`reduce`" operation does the reduction in dimension, by yielding the sum of all elements in the array as a

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

scalar (sum), while the prefix sum yields an array of the same size, with each element containing the "prefix-sum" for that index.

Workgroup broadcast built-ins are also used.

The following code illustrates how the workgroup built-ins in OpenCL 2.0 can be used to efficiently extract the primes from an array of positive integers.

```
int isPrime(int number) {
    int i;
    if (number == 0) return 0;
    for (i=2; i*i<=number; i++) {
        if (number % i == 0) return 0;
    }
    return 1;
}

/***
 * set_primes_kernel:
 * this kernel fills the boolean primes array to reflect the entry
in
 * input array is prime
***/
kernel void set_primes_kernel(global int *in, global int *primes)
{
  int id = get_global_id(0);

  primes[id] = 0;
  if (isPrime(in[id]))
        primes[id] = 1;

}

/***
 * get_primes_kernel:
 * this kernel fills output array to reflect the entry in
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
 * input array is prime

***/

kernel void get_primes_kernel(global int *in, global int *output,
global int *outPrimes)

{

  int id = get_global_id(0);

  int k = output[id] - 1;


  if (id == 0) {

   if (output[id] == 1)

          outPrimes[k] = in[id];

   return;

   }


  if (output[id-1] == k)

       outPrimes[k] = in[id];


}


/***

 * group_scan_kernel:

 * this kernel takes an input array [in] and produces a work-group
level scan in [out].

 ***/

kernel void group_scan_kernel(global int *in, global int *out)

{

  int  in_data;


  int i = get_global_id(0);


  in_data = in[i];

  out[i]  = work_group_scan_inclusive_add(in_data);

}
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
/***
 * global_scan_kernel:
 * takes a work-group scanned array [out] from group_scan_kernel
and gives globally scanned array.
 ***/
kernel void global_scan_kernel(global int *out, unsigned int stage)
{
  unsigned int l;
  unsigned int vlen;

  unsigned int prev_gr,prev_el;
  unsigned int curr_gr, curr_el;

  int    add_elem;

  int lid  = get_local_id(0);
  int grid = get_group_id(0);
  int szgr = get_local_size(0);

  /* array size to be processed */
  vlen = 1 << stage;

  /* find the element to be added */
  l       = (grid >> stage);
  prev_gr = l*(vlen << 1) + vlen - 1;
  prev_el = prev_gr*szgr + szgr - 1;

  if (lid == 0)
    add_elem    = out[prev_el];

  work_group_barrier(CLK_GLOBAL_MEM_FENCE|CLK_LOCAL_MEM_FENCE);
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
  add_elem = work_group_broadcast(add_elem,0);


  /* find the array to which the element to be added */
  curr_gr = prev_gr + 1 + (grid % vlen);
  curr_el = curr_gr*szgr + lid;
  out[curr_el] += add_elem;
}
```

The first kernel fills the "primes" Boolean array as in the OpenCL 1.2 example. The second kernel runs "group_scan_kernel" on this boolean "prime" array. However, the "group_scan_kernel" function works only at the workgroup level. The next kernel, "global_scan_kernel", runs in multiple (`log (n)`) stages, where "n" is the size of the input array. At each stage, it merges two consecutive work_groups for which the "prefix_sums" are ready for the work group. After log(n) stages, the prefix sums for the complete "out" array are computed. This gives index of the final primes array. During every stage, work_group_broadcast is used to broadcast the prefix sum of the last element of the previous workgroup. The following figures show the outputs of various arrays at various stages.

Input:

12	31	47	64	19	27	49	81	99	11

(Assuming workgroup size is 4)

After setPrimes (Primes):

0	1	1	0	1	0	0	0	0	1

After group_scan_kernel (out):

0	1	2	2	1	1	1	1	0	1

After global_scan_kernel out:

0	1	2	2	3	3	3	3	3	4

Finally, the device enqueue feature is employed. Device enqueue enqueues these kernels in sequence instead of host-enqueing them. This improves performance. As these kernels must be run in sequence, they cannot be clubbed into one kernel.

outPrimes:

31	47	19	11

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

The host launches the setPrimes, group_scan, and global_scan kernels in sequence to set the primes Boolean array, performing prefix sum calculation at the workgroup level, and perform the prefix sum across workgroup levels, respectively. The output after each stage is shown in the above figure.

The performance of this OpenCL 2.0 version is compared with that of OpenCL 1.2 version for the same problem. Unlike OpenCL 2.0, OpenCL 1.2 does not provide workgroup built-ins to perform prefix sum efficiently. While primes can be in the same manner as in OpenCL 2.0, the group and global scan contain redundant computations of the order of O(n) for each work item, as illustrated in the following basic OpenCL 1.2 implementation:

```
kernel void get_primes_kernel(global int *in, global int *primes,
global int *outPrimes)

{

  int id = get_global_id(0);

  int idx,i;


  if (primes[id]) {

       idx = 0;

       for (i=0;i<id;i++)

      idx += primes[i];


       outPrimes[idx] = in[id];

       }


}
```

Thus, in OpenCL 1.2, each work item computes the prefix sum for its index and stores the input prime into the outPrimes array. Clearly, more work is performed for each work item, compared to OpenCL 2.0.

The following graph and the table compare the performance of OpenCL 2.0 version with that of the OpenCL 1.2 version:

Size	OpenCL 1.2 (Secs)	OpenCL 2.0 (Secs)
1024	0.000365	0.000495
2048	0.000588	0.0006
4096	0.00115	0.0005
10000	0.003	0.0011
20000	0.00692	0.0012

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

40000	0.02222	0.00165
80000	0.07897	0.00254
160000	0.29292	0.00391
320000	1.1427	0.00614
640000	4.5203	0.010622
128000	17.991	0.01928

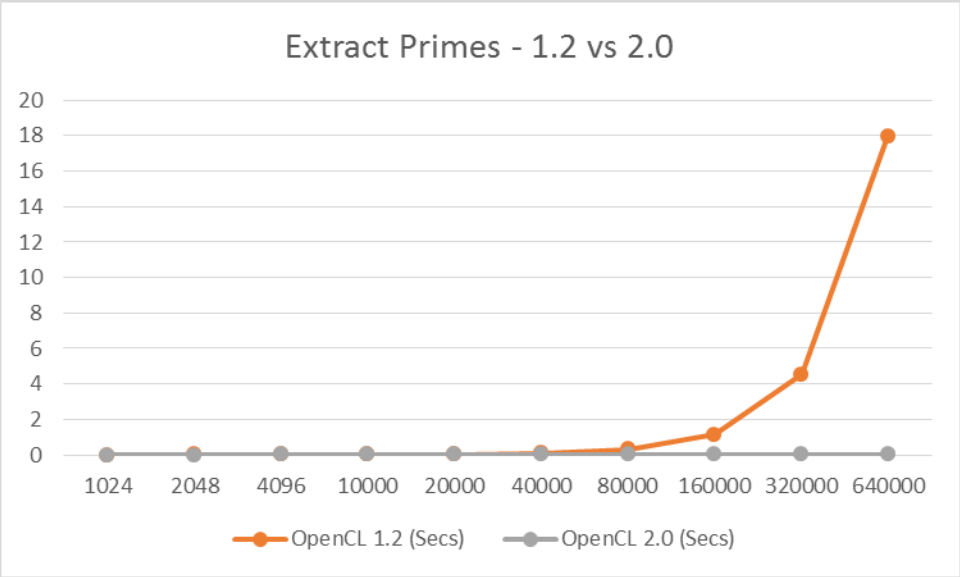

As the size of the input array exceeds 4K, the difference in performance between the OpenCL 1.2 implementation and the OpenCL 2.0 implementation becomes more apparent. In OpenCL 1.2, the time taken is O(n) time and the work done is O(n^2); in OpenCL 2.0, the time taken is O(log (n)) and the work done is O(n), assuming the prefix sum at the group level is constant time.

6.4.2.4 Binary search using device enqueue (or kernel enqueue)

The power of device enqueue is aptly illustrated in the example of binary search. To make the problem interesting, multiple keys in a sorted array will be searched for. The versions written for OpenCL 1.2 and 2.0 will also be compared with respect to programmability and performance.

The OpenCL 1.2 version of the code that performs binary search is as follows:

```
kernel void binarySearch_mulkeys( global int *keys, global uint
*input, const unsigned int numKeys, global int *output)

{

    int gid = get_global_id(0);

    int lBound = gid * 256;
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
    int uBound = lBound + 255;
    for(int i = 0; i < numKeys; i++)
    {
        if(keys[i] >= input[lBound] && keys[i] <=input[uBound])
            output[i]=lBound;
     }
}
```

The search for multiple keys is done sequentially, while the sorted array is divided into 256 sized chunks. The NDRange is the size of the array divided by the chunk size. Each work item checks whether the key is present in the range and if the key is present, updates the output array.

The issue with the above approach is that if the input array is very large, the number of work items (NDRange) would be very large. The array is not divided into smaller, more-manageable chunks.

In OpenCL 2.0, the device enqueue feature offers clear advantages in binary search performance.

The kernel is rewritten in OpenCL 2.0 as follows:

```
for(key_count=0;key_count<no_of_keys;key_count++)
   {
          /* If the element to be found does not lie between them,
then nothing left to do in this thread */
          if( (elementLower > keys[key_count]) || (elementUpper <
keys[key_count]))
                  continue;
          else
          {
                /* However, if the element does lie between the
lower and upper bounds of this thread's searchspace
                 * we need to narrow down the search further in
this search space
                 */

            subdivSize_for_keys[key_count] = subdivSize;
            outputArray[key_count].x = parent_globalids[key_count]
+ tid * subdivSize_for_keys[key_count];
            parent_globalids[key_count]  +=
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
subdivSize_for_keys[key_count] * tid ;

            outputArray[key_count].w = 1;

            outputArray[key_count].y =
subdivSize_for_keys[key_count];

            globalLowerIndex = tid*subdivSize_for_keys[key_count];

            unsigned int globalUpperIndex =
globalLowerIndex+subdivSize_for_keys[key_count]-1;

            if(subdivSize_for_keys[key_count] < global_threads)

                 continue;

            subdivSize_for_keys[key_count] = (globalUpperIndex -
globalLowerIndex + 1)/global_threads;

            ndrange_t ndrange1 = ndrange_1D(global_threads);

           void (^binarySearch_device_enqueue_wrapper_blk)(void) =
^{binarySearch_device_enqueue(outputArray,

                        sortedArray,

                        subdivSize_for_keys[key_count],

                        globalLowerIndex,


keys[key_count],global_threads,key_count,parent_globalids);};



                        int err_ret =
enqueue_kernel(def_q,CLK_ENQUEUE_FLAGS_NO_WAIT,ndrange1,binarySearc
h_device_enqueue_wrapper_blk);

                        if(err_ret != 0)

                        {

                               outputArray[key_count].w = 2;

                               outputArray[key_count].z = err_ret;

                              return;

                        }

          }
```

In the OpenCL 2.0 version, each work item checks for each key, if it is found in its search range. If the key is found, it further divides the range into chunks and enqueues the kernel for further processing.

The advantage is that when the input array is large, the OpenCL 2.0 version divides the input array into 1024-sized chunks. The chunk in which the given key falls is found and another kernel is enqueued which further divides it into 1024-

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

sized chunks, and so on. In OpenCL 1.2, as the whole array is taken as the NDRange, a huge number of work groups require processing.

The following figure shows how the OpenCL 2.0 version compares to the OpenCL 1.2 as the array increases beyond a certain size.

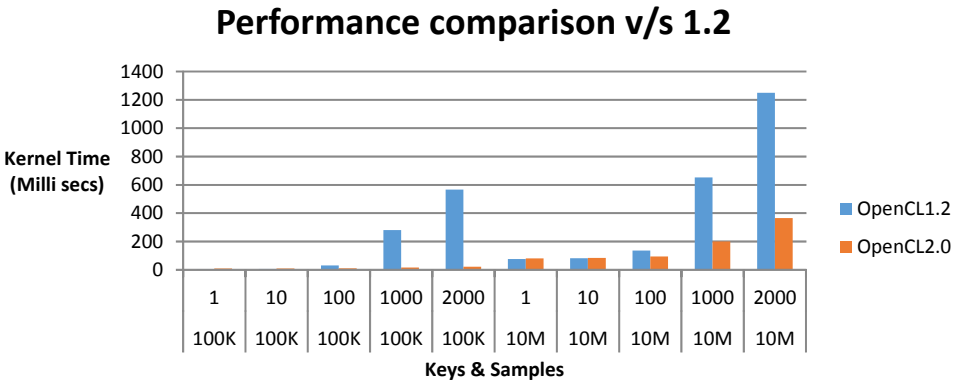

The above figure shows the performance benefit of using OpenCL 2.0 over the same sample using OpenCL 1.2. In OpenCL 2.0, the reduced number of kernel launches from the host allow superior performance. The kernel enqueues are much more efficient when done from the device.

6.5 Atomics and synchronization

6.5.1 Overview

In OpenCL 1.2, only work-items in the same workgroup can synchronize.

OpenCL 2.0 introduces a new and detailed memory model which allows developers to reason about the effects of their code on memory, and in particular understand whether atomic operations and fences used for synchronization ensure the visibility of variables being used to communicate between threads. In conjunction with the new memory model, OpenCL 2.0 adds a new set of atomic built-in functions and fences derived from C++11 (although the set of types is restricted), and also deprecates the 1.2 atomic built in functions and fences.

These additions allow synchronization between work-items in different work-groups, as well as fine-grained synchronization with the host using atomic operations on memory in fine-grained SVM buffers (allocated with the `CL_MEM_SVM_ATOMICS` flag).

6.5.2 Usage

The following examples to illustrate the use of atomics are part of the AMD APP SDK.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

6.5.2.1 Atomic Loads/Stores

This sample illustrates atomic loads/stores with the use of memory orders.

The first step is to create this memory on the host:

```
buffer = (int *) clSVMAllocP(context, CL_MEM_SVM_FINE_GRAIN_BUFFER,
(N+1)*sizeof(int), 4);

atomicBuffer = (int *) clSVMAllocP(context,
CL_MEM_SVM_FINE_GRAIN_BUFFER | CL_MEM_SVM_ATOMICS,
(N+1)*sizeof(int), 4);
```

Note the flags sent as parameters: `CL_MEM_SVM_FINE_GRAIN_BUFFER` and `CL_MEM_SVM_ATOMICS`. The following kernel runs on all work items in parallel. It will atomically load `atomicBuffer[0]`, check whether its value is 99, and wait till it is 99. The acquire memory order is used to indicate that the latest update must be done on the host and to ensure that the local L1 cache is not read from. This will be made 99 by the host (CPU) by

```
std::atomic_store_explicit ((std::atomic<int> *)&atomicBuffer[0],
99, std::memory_order_release);
```

The host uses the C++11 compiler and the same memory model.

```
kernel void ldstore(volatile global int *buffer, global int*
atomicBuffer)

{

        int i;

        while (atomic_load_explicit ((global atomic_int
*)&atomicBuffer[0], memory_order_acquire) != 99);

        i = get_global_id(0);

         buffer[i] += i;

         atomic_store_explicit ((global atomic_int
*)&atomicBuffer[i], (100+i), memory_order_release);

}
```

The kernel next stores `(100+i)`, where `i` is the ID of the work-item into `atomicBuffer[i]`. The order used is `memory_order_release` which ensures that the updated copy reaches the CPU which is waiting for it to report `PASS` for the test.

After the atomic operation, the updates on fine-grain variables (such as buffer) will also be available at the host. The CPU checks for the following to ensure that the results are OK:

```
for  (i=0;i<N;i++)
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
   while(std::atomic_load_explicit ((std::atomic<int>
*)&atomicBuffer[i], std::memory_order_acquire) != (100+i));

/* check the results now */

for (i=0;i<N;i++)

if (buffer[i] != (64+i))

     printf(" Test Failed \n");

printf (" Test Passed! \n");
```

6.5.2.2 Atomic Compare and Exchange (CAS)

This sample illustrates the use of the atomic CAS operation typically used for "lock-free" programming, in which a critical section can be created without having to use waiting mutexes/semaphores. The following kernel simultaneously inserts the IDs of various work items into the "list" array by using atomic CAS operation. The same loop also runs on the host and inserts the other half (N) work items. In this way, 2*N numbers are inserted into this "list".

```
kernel void linkKernel(__global int *list) {

int head, i;

i = get_global_id(0) + 1;

head = list[0];

if (i != get_global_size(0)) {

   do {

          list[i] = head;

         } while (!atomic_compare_exchange_strong((global
atomic_int *) &list[0], &head,i), memory_order_release,
memory_order_acquire, memory_scope_system);

   }

}
```

Note how there is no wait to enter the critical section, but list[0] and head are updated atomically. On the CPU too, a similar loop runs. Again note that the variables "list"and "head" must be in fine-grain SVM buffers. `memory_order_release` and `memory_scope_system` are used to ensure that the CPU gets the updates -- hence the name "platform atomics."

6.5.2.3 Atomic Fetch

This sample illustrates the use of the atomic fetch operation. The fetch operation is an RMW (Read-Modify-Write) operation. The following kernel computes the maximum of the N numbers in array "A". The result of the intermediate comparisons is computed and the result is placed in a Boolean array "B". After

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

the matrix "B" is computed, the row (i) is computed. The row which has all 1s will be the maximum (C[i]).

```
kernel void atomicMax(volatile global int *A, global int *B, global
int *C, global int *P)

{

int  i = get_global_id(0);

   int  j = get_global_id(1);

   int N = *P, k;

   if (A[i] >= A[j]) B[i*N+j] = 1;

   else B[i*N+j] = 0;

   if (j == 0) {

          C[i] = 1;

           for (k=0;k<N;k++)

            atomic_fetch_and_explicit((global atomic_int *)&C[i],
B[i*N+k], memory_order_release, memory_scope_device);

   }

}
```

Similarly, another sample includes the following kernel that increments 2*N times, N times in the kernel and another N times on the host:

```
kernel void counter(__global int *count)

{

   atomic_fetch_add((atomic _int)count, 1);

   //(*count)++;

}
```

Note: If `atomic_fetch_add` is not used and instead an incrementing count (as performed in the commented line) is used, the sum will not be computed correctly.

6.6 Pipes

6.6.1 Overview

OpenCL 2.0 introduces a new mechanism, pipes, for passing data between kernels. A pipe is essentially a structured buffer containing space for some number of a single kernel specified type of "packet" and bookkeeping information. Pipes are accessed via special `read_pipe` and `write_pipe` built-in functions. A given kernel may either read from or write to a pipe, but not both. Pipes are only "coherent" at the standard synchronization points; the result of

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

concurrent accesses to the same pipe by multiple kernels (even if permitted by hardware) is undefined. A pipe cannot be accessed from the host side.

Pipes are created on the host with a call to `clCreatePipe`, and may be passed between kernels. Pipes may be particularly useful when combined with device-size enqueue for dynamically constructing computational data flow graphs.

There are two types of pipes: a read pipe, from which a number of packets can be read; and a write pipe, to which a number of packets can be written.

Note: A pipe specified as read-only cannot be written into and a pipe specified as write-only cannot be read from. A pipe cannot be read from and written into at the same time.

6.6.2 Functions for accessing pipes

A new host API function has been added into the OpenCL 2.0 spec to create the Pipe.

```
cl_mem    clCreatePipe (  cl_context context,
cl_mem_flags flags,
cl_uint pipe_packet_size,
cl_uint pipe_max_packets,
const cl_pipe_properties * properties,
cl_int *errcode_ret)
```

The memory allocated in the above function can be passed to kernels as read-only or write-only pipes.

Also, a set of built-in functions have been added to operate on the pipes.

Read pipe and write pipe functions	
int read_pipe (pipe gentype p, gentype *ptr)	Read packet from pipe p into ptr. Returns 0 if read_pipe is successful and a negative value otherwise.
int write_pipe (pipe gentype p,const gentype *ptr)	Write packet specified by ptr to pipe p. Returns 0 if write_pipe is successful and a negative value otherwise.
Pipe Reservation functions	
reserve_id_t reserve_read_pipe (pipe gentype p,uint num_packets) reserve_id_t reserve_write_pipe (pipe gentype p,uint num_packets) reserve_id_t work_group_reserve_read_pipe (pipe gentype p,uint num_packets) reserve_id_t work_group_reserve_write_pipe (pipe gentype p,uint num_packets) reserve_id_t sub_group_reserve_read_pipe (pipe gentype p,uint num_packets) reserve_id_t sub_group_reserve_write_pipe (pipe gentype p,uint num_packets)	Reserve num_packets entries for reading from or writing to pipe p. Returns a valid reservation ID if the reservation is successful. The reserved pipe entries are referred to by indices that go from 0 … num_packets – 1.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

bool is_valid_reserve_id (reserve_id_t reserve_id)	Return true if reserve_id is a valid reservation ID and false otherwise.
Read and write pipe functions using reservation indices	
int read_pipe (pipe gentype p, reserve_id_t reserve_id, uint index, gentype *ptr)	Read packet from the reserved area of the pipe referred to by reserve_id and index into ptr. The reserved pipe entries are referred to by indices that go from 0 … num_packets – 1. Returns 0 if read_pipe is successful and a negative value otherwise.
int write_pipe (pipe gentype p, reserve_id_t reserve_id, uint index, const gentype *ptr)	Write packet specified by ptr to the reserved area of the pipe referred to by reserve_id and index. The reserved pipe entries are referred to by indices that go from 0 … num_packets – 1. Returns 0 if write_pipe is successful and a negative value otherwise.
Commit/flush functions	
void commit_read_pipe (pipe gentype p, reserve_id_t reserve_id) void commit_write_pipe (pipe gentype p, reserve_id_t reserve_id) void work_group_commit_read_pipe (pipe gentype p, reserve_id_t reserve_id) void work_group_commit_write_pipe (pipe gentype p, reserve_id_t reserve_id) void sub_group_commit_read_pipe (pipe gentype p, reserve_id_t reserve_id) void sub_group_commit_write_pipe (pipe gentype p, reserve_id_t reserve_id)	Indicates that all reads and writes to num_packets associated with reservation reserve_id are completed.
Pipe Query Functions	
uint get_pipe_num_packets (pipe gentype p)	Returns the number of available entries in the pipe. The number of available entries in a pipe is a dynamic value. The value returned should be considered immediately stale.
uint get_pipe_max_packets (pipe gentype p)	Returns the maximum number of packets specified when pipe was created.

6.6.3 Usage

The corresponding sample in the AMD APP SDK contains two kernels: a "producer_kernel", which writes into the pipe, and another kernel, "consumer_kernel", which reads from the pipe. In this example, the producer writes a sequence of random numbers. The consumer reads these numbers and creates a histogram of these numbers. The producer kernel first reserves the amount of space in the pipe by invoking:

```
rid = work_group_reserve_write_pipe(rng_pipe, szgr);
```

Next, it writes and commits into the pipe by invoking:

```
write_pipe(rng_pipe,rid,lid, &gfrn);

work_group_commit_write_pipe(rng_pipe, rid);
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Similarly, the consumer pipe reserves the pipe for reading and reads from the pipe by issuing:

```
work_group_reserve_read_pipe(rng_pipe, szgr);

read_pipe(rng_pipe,rid,lid, &rn);

work_group_commit_read_pipe(rng_pipe, rid);
```

The pipe thus provides a useful communication mechanism between two kernels.

6.7 Sub-groups

6.7.1 Overview

OpenCL 2.0 introduces a Khronos sub-group extension. Sub-groups are a logical abstraction of the hardware SIMD execution model akin to wavefronts, warps, or vectors and permit programming closer to the hardware in a vendor-independent manner. This extension includes a set of cross-sub-group built-in functions that match the set of the cross-work-group built-in functions specified in Section 6.4

6.8 Program-scope global Variables

6.8.1 Overview

OpenCL 1.2 permits the declaration of only constant address space variables at program scope.

OpenCL 2.0 permits the declaration of variables in the global address space at program (i.e. outside function) scope. These variables have the lifetime of the program in which they appear, and may be initialized. The host cannot directly access program-scope variables; a kernel must be used to read/write their contents from/to a buffer created on the host.

Program-scope global variables can save data across kernel executions. Using program-scope variables can potentially eliminate the need to create buffers on the host and pass them into each kernel for processing. However, there is a limit to the size of such variables. The developer must ensure that the total size does not exceed the value returned by the device info query: `CL_DEVICE_MAX_GLOBAL_VARIABLE_SIZE`.

6.9 Image Enhancements

6.9.1 Overview

OpenCL 2.0 introduces significant enhancements for processing images.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

A read_write access qualifier for images has been added. The qualifier allows reading from and writing to certain types of images (verified against clGetSupportedImageFormats by using the CL_MEM_KERNEL_READ_AND_WRITE flag) in the same kernel, but reads must be sampler-less. An atomic_work_item_fence with the CLK_IMAGE_MEM_FENCE flage and the memory_scope_work_item memory scope is required between reads and writes to the same image to ensure that the writes are visible to subsequent reads. If multiple work-items are writing to and reading from multiple locations in an image, a call to work_group_barrier with the CLK_IMAGE_MEM_FENCE flag is required.

OpenCL 2.0 also allows 2D images to be created from a buffer or another 2D image and makes the ability to write to 3D images a core feature. This extends the power of image operations to more situations.

The function clGetSupportedImageFormats returns a list of the image formats supported by the OpenCL platform. The Image format has two parameters, channel order and data type. The following lists some image formats OpenCL supports:

Channel orders: CL_A, CL_RG, CL_RGB, CL_RGBA.
Channel data type: CL_UNORM_INT8, CL_FLOAT.

OpenCL 2.0 provides improved image support, specially support for sRGB images and depth images.

6.9.2 sRGB

sRGB is a standard RGB color space that is used widely on monitors, printers, digital cameras, and the Internet. Because the linear RGB value is used in most image processing algorithms. Processing the images often requires converting sRGB to linear RGB.

OpenCL 2.0 provides a new feature for handling this conversion directly. Note that only the combination of data type CL_UNORM_INT8 and channel order CL_sRGBA is mandatory in OpenCL 2.0. The AMD implementations support this combination. The remaining combinations are optional in OpenCL 2.0.

When not using the mandatory combination (CL_sRGBA, CL_UNORM_INT8), the clGetSupportedImageFormats function must be used to get a list of supported image formats and data types before using the sRGB image,

Creating sRGB image objects is similar to creating an image object of existing supported channel order with OpenCL 2.0. The following snippet shows how to create CL_sRGBA image objects by using the read_image call.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
cl_image_format  imageFormat;
imageFormat.image_channel_data_type  = CL_UNORM_INT8;
imageFormat.image_channel_order = CL_sRGBA
cl_mem imageObj = clCreateImage(
context,                          // A valid OpenCL context
CL_MEM_READ_ONY | CL_MEM_COPY_HOST_PTR,
&imageFormat,
&desc, //cl_image_desc
pSrcImage,                    // An pointer to the image data
&retErr);                       // Returned error code
```

A new sRGB image can also be created based on an existing RGB image object, so that the kernel can implicitly convert the sRGB image data to RGB. This is useful when the viewing pixels are sRGB but share the same data as the existing RGB image.

After an sRGB image object has been created, the `read_imagef` call can be used in the kernel to read it transparently. `read_imagef` explicitly converts sRGB values into linear RGB. Converting sRGB into RGB in the kernel explicitly is not necessary if the device supports OpenCL 2.0. Note that only `read_imagef` can be used for reading sRGB image data because only the `CL_UNOR_INT8` data type is supported with OpenCL 2.0.

The following is a kernel sample that illustrates how to read an sRGB image object.

```
// Read sRGBA image object (input) and convert it to linear RGB
values(results)
__kernel void sample_kernel( read_only image2d_t input, sampler_t
imageSampler,  __global float *xOffsets, __global float *yOffsets,
__global float4 *results  )   // input: sRGBA image object
{
     int tidX = get_global_id(0), tidY = get_global_id(1);
     int offset = tidY*get_image_width(input) + tidX;
     int2 coords = (int2)( xOffsets[offset], yOffsets[offset]);
     results[offset] = read_imagef( input, imageSampler, coords  );
}
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

OpenCL 2.0 does not include writing sRGB images directly, but provides the `cl_khr_srgb_image_writes` extension. The AMD implementations do not support this extension as of this writing.

In order to write sRGB pixels in a kernel, explicit conversion from linear RGB to sRGB must be implemented in the kernel.

`clFillImage` is an exception for writing sRGB image directly. The AMD OpenCL platform supports `clFillImage` for filling linear RGB image to sRGB image directly.

6.9.3 Depth images

As with other image formats, `clCreateImage` is used for creating depth image objects. However, the channel order must be set to `CL_DEPTH`, as illustrated below. For the data type of depth image, OpenCL 2.0 supports only `CL_FLOAT` and `CL_UNORM_INT16`.

```
cl_image_format  imageFormat;

imageFormat.image_channel_data_type  = CL_UNORM_INT16;

imageFormat.image_channel_order = CL_DEPTH

cl_mem imageObj = clCreateImage(

                            context,                            // A
valid OpenCL context

                           CL_MEM_READ_ONY | CL_MEM_COPY_HOST_PTR,

                            &imageFormat,

                            &desc,
//cl_image_desc

                            pSrcImage,                      // An
pointer to the image data

                            &retErr);                         //
Returned error code
```

In OpenCL 2.0, depth images must be of type `image2d` or `image2d array`. `clCreateImage` will fail for other dimensions when creating depth image.

A depth image object can be read by using the `read_imagef` call in the kernel. For write, `write_imagef` must be used. `read_image(i|ui)` and `write_image(i|ui)` are not supported for depth images.

OpenCL 2.0 C introduces two data types, `image2d_depth_t` and `image2d_array_depth_t` for declaring depth images. The following kernel code sample illustrates how to read depth image objects.

```
// Read depth image object (input) based on sampler and offset and
save it (results)

__kernel void sample_kernel( read_only image2d_depth_t input,
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
sampler_t imageSampler, __global float *xOffsets, __global float
*yOffsets, __global float *results  )

{

   int tidX = get_global_id(0), tidY = get_global_id(1);

   int offset = tidY*get_image_width(input) + tidX;

   int2 coords = (int2)( xOffsets[offset], yOffsets[offset]);

   results[offset] = read_imagef( input, imageSampler, coords  );

}
```

The AMD OpenCL 2.0 platform fully supports the `cl_khr_depth_images` extension but not the `cl_khr_gl_depth_images` extension. Consequently, the AMD OpenCL platform does not support creating a CL depth image from a GL depth or depth-stencil texture.

6.10 Non-uniform work group size

6.10.1 Overview

Prior to OpenCL 2.0, each work-group size needed to divide evenly into the corresponding global size. This requirement is relaxed in OpenCL 2.0; the final work-group in each dimension is allowed to be smaller than all of the other work-groups in the "uniform" part of the NDRange. This can reduce the effort required to map problems onto NDRanges.

A consequence is that kernels may no longer assume that calls to `get_work_group_size` return the same value in all work-groups. However, a new call (`get_enqueued_local_size`) has been added to obtain the size in the uniform part, which is specified using the `local_work_size` argument to the `clEnqueueNDRangeKernel`.

A new compile time option (`-cl-uniform-work-group-size`) has been added to optimize the computation for cases in which the work-group size is known to, or required to, divide evenly into the global size.

6.11 Portability considerations

6.11.1 Migrating from OpenCL 1.2 to OpenCL 2.0

OpenCL 2.0 is backward compatible with OpenCL 1.2. Applications written on OpenCL 1.2 should run on OpenCL 2.0 without requiring any changes to the application.

OpenCL 2.0 includes changes in the runtime and the compiler. In the runtime, some new functions (such as for SVM) have been added. In the compiler, the `-cl-std=CL2.0` option is needed in order to compile OpenCL 2.0 kernels.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

If a program uses the OpenCL 2.0 functions and if one compiles a kernel by using the `cl-std=CL2.0` option, the program will not build or compile on OpenCL 1.2 platforms. If a program uses only OpenCL 1.2 functions and if one compiles a kernel without the `cl-std=CL2.0` option, then the program should run on OpenCL 2.0 platforms.

6.11.2 Identifying implementation specifics

Applications can query for the OpenCL extensions and use the values returned from the OpenCL functions.

For instance, `clGetSupportedImageFormats` will return all image formats supported by OpenCL. The supported images may differ across implementations. Similarly, `clGetDeviceInfo` with the `CL_DEVICE_EXTENSIONS` parameter returns all the supported extensions. The supported extensions may differ across implementations and between different versions of OpenCL.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Appendix A OpenCL Optional Extensions

The OpenCL extensions are associated with the devices and can be queried for a specific device. Extensions can be queried for platforms also, but that means that all devices in the platform support those extensions.

Table A.1, on page A-15, lists the supported extensions.

A.1 Extension Name Convention

The name of extension is standardized and must contain the following elements without spaces in the name (in lower case):

- `cl_khr_<extension_name>` - for extensions approved by Khronos Group. For example: `cl_khr_fp64`.
- `cl_ext_<extension_name>` - for extensions provided collectively by multiple vendors. For example: `cl_ext_device_fission`.
- `cl_<vendor_name>_<extension_name>` – for extension provided by a specific vendor. For example: `cl_amd_media_ops`.

The OpenCL Specification states that all API functions of the extension must have names in the form of `cl<FunctionName>KHR`, `cl<FunctionName>EXT`, or `cl<FunctionName><VendorName>`. All enumerated values must be in the form of `CL_<enum_name>_KHR`, `CL_<enum_name>_EXT`, or `CL_<enum_name>_<VendorName>`.

A.2 Querying Extensions for a Platform

To query supported extensions for the OpenCL platform, use the `clGetPlatformInfo()` function, with the `param_name` parameter set to the enumerated value `CL_PLATFORM_EXTENSIONS`. This returns the extensions as a character string with extension names separated by spaces. To find out if a specific extension is supported by this platform, search the returned string for the required substring.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

A.3 Querying Extensions for a Device

To get the list of devices to be queried for supported extensions, use one of the following:

- Query for available platforms using `clGetPlatformIDs()`. Select one, and query for a list of available devices with `clGetDeviceIDs()`.
- For a specific device type, call `clCreateContextFromType()`, and query a list of devices by calling `clGetContextInfo()` with the param_name parameter set to the enumerated value `CL_CONTEXT_DEVICES`.

After the device list is retrieved, the extensions supported by each device can be queried with function call `clGetDeviceInfo()` with parameter `param_name` being set to enumerated value `CL_DEVICE_EXTENSIONS`.

The extensions are returned in a char string, with extension names separated by a space. To see if an extension is present, search the string for a specified substring.

A.4 Using Extensions in Kernel Programs

There are special directives for the OpenCL compiler to enable or disable available extensions supported by the OpenCL implementation, and, specifically, by the OpenCL compiler. The directive is defined as follows.

```
#pragma OPENCL EXTENSION <extention_name> : <behavior>
#pragma OPENCL EXTENSION all: <behavior>
```

The *<extension_name>* is described in Section A.1, “Extension Name Convention.”. The second form allows to address all extensions at once.

The *<behavior>* token can be either:

- `enable` - the extension is enabled if it is supported, or the error is reported if the specified extension is not supported or token “all” is used.
- `disable` - the OpenCL implementation/compiler behaves as if the specified extension does not exist.
- `all` - only core functionality of OpenCL is used and supported, all extensions are ignored. If the specified extension is not supported then a warning is issued by the compiler.

The order of directives in `#pragma OPENCL EXTENSION` is important: a later directive with the same extension name overrides any previous one.

The initial state of the compiler is set to ignore all extensions as if it was explicitly set with the following directive:

```
#pragma OPENCL EXTENSION all : disable
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

This means that the extensions must be explicitly enabled to be used in kernel programs.

Each extension that affects kernel code compilation must add a defined macro with the name of the extension. This allows the kernel code to be compiled differently, depending on whether the extension is supported and enabled, or not. For example, for extension `cl_khr_fp64` there should be a `#define` directive for macro `cl_khr_fp64`, so that the following code can be preprocessed:

```
#ifdef cl_khr_fp64
    // some code
#else
    // some code
#endif
```

A.5 Getting Extension Function Pointers

Use the following function to get an extension function pointer.

`void* clGetExtensionFunctionAddress(const char* FunctionName).`

This returns the address of the extension function specified by the `FunctionName` string. The returned value must be appropriately cast to a function pointer type, specified in the extension spec and header file.

A return value of NULL means that the specified function does not exist in the CL implementation. A non-NULL return value does not guarantee that the extension function actually exists – queries described in sec. 2 or 3 must be done to ensure the extension is supported.

The `clGetExtensionFunctionAddress()` function cannot be used to get core API function addresses.

A.6 List of Supported Extensions that are Khronos-Approved

Supported extensions approved by the Khronos Group are:

- `cl_khr_global_int32_base_atomics` – basic atomic operations on 32-bit integers in global memory.
- `cl_khr_global_int32_extended_atomics` – extended atomic operations on 32-bit integers in global memory.
- `cl_khr_local_int32_base_atomics` – basic atomic operations on 32-bit integers in local memory.
- `cl_khr_local_int32_extended_atomics` – extended atomic operations on 32-bit integers in local memory.
- `cl_khr_int64_base_atomics` – basic atomic operations on 64-bit integers in both global and local memory.
- `cl_khr_int64_extended_atomics` – extended atomic operations on 64-bit integers in both global and local memory.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

- cl_khr_3d_image_writes – supports kernel writes to 3D images.
- cl_khr_byte_addressable_store – this eliminates the restriction of not allowing writes to a pointer (or array elements) of types less than 32-bit wide in kernel program.
- cl_khr_gl_sharing – allows association of OpenGL context or share group with CL context for interoperability.
- cl_khr_icd – the OpenCL Installable Client Driver (ICD) that lets developers select from multiple OpenCL runtimes which may be installed on a system. This extension is automatically enabled as of SDK v2 for AMD Accelerated Parallel Processing.
- cl_khr_d3d10_sharing - allows association of D3D10 context or share group with CL context for interoperability.

A.7 cl_ext Extensions

- cl_ext_device_fission - Support for device fission in OpenCL™. For more information about this extension, see:

  http://www.khronos.org/registry/cl/extensions/ext/cl_ext_device_fission.txt

A.8 AMD Vendor-Specific Extensions

This section describes the AMD vendor-specific extensions.

A.8.1 cl_amd_fp64

Before using double data types, double-precision floating point operators, and/or double-precision floating point routines in OpenCL™ C kernels, include the `#pragma OPENCL EXTENSION cl_amd_fp64 : enable` directive. See Table A.1 for a list of supported routines.

A.8.2 cl_amd_vec3

This extension adds support for vectors with three elements: float3, short3, char3, etc. This data type was added to OpenCL 1.1 as a core feature. For more details, see section 6.1.2 in the OpenCL 1.1 or OpenCL 1.2 spec.

A.8.3 cl_amd_device_persistent_memory

This extension adds support for the new buffer and image creation flag CL_MEM_USE_PERSISTENT_MEM_AMD. Buffers and images allocated with this flag reside in host-visible device memory. This flag is mutually exclusive with the flags CL_MEM_ALLOC_HOST_PTR and CL_MEM_USE_HOST_PTR.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

A.8.4 cl_amd_device_attribute_query

This extension provides a means to query AMD-specific device attributes. To enable this extension, include the `#pragma OPENCL EXTENSION cl_amd_device_attribute_query : enable` directive. Once the extension is enabled, and the `clGetDeviceInfo` parameter `<param_name>` is set to `CL_DEVICE_PROFILING_TIMER_OFFSET_AMD`, the offset in nano-seconds between an event timestamp and Epoch is returned.

1.8.4.1 cl_device_profiling_timer_offset_amd

This query enables the developer to get the offset between event timestamps in nano-seconds. To use it, compile the kernels with the `#pragma OPENCL EXTENSION cl_amd_device_attribute_query : enable directive`. For kernels complied with this pragma, calling `clGetDeviceInfo` with <*param_name*> set to `CL_DEVICE_PROFILING_TIMER_OFFSET_AMD` returns the offset in nano-seconds between event timestamps.

1.8.4.2 cl_amd_device_topology

This query enables the developer to get a description of the topology used to connect the device to the host. Currently, this query works only in Linux. Calling `clGetDeviceInfo` with <*param_name*> set to `CL_DEVICE_TOPOLOGY_AMD` returns the following 32-bytes union of structures.

```
typedef union
{
    struct { cl_uint type; cl_uint data[5]; } raw;
    struct { cl_uint type; cl_char unused[17]; cl_char bus; cl_char
device; cl_char function; } pcie; } cl_device_topology_amd;
```

The type of the structure returned can be queried by reading the first unsigned int of the returned data. The developer can use this type to cast the returned union into the right structure type.

Currently, the only supported type in the structure above is PCIe (type value = 1). The information returned contains the PCI Bus/Device/Function of the device, and is similar to the result of the `lspci` command in Linux. It enables the developer to match between the OpenCL device ID and the physical PCI connection of the card.

1.8.4.3 cl_amd_device_board_name

This query enables the developer to get the name of the GPU board and model of the specific device. Currently, this is only for GPU devices.

Calling `clGetDeviceInfo` with <*param_name*> set to `CL_DEVICE_BOARD_NAME_AMD` returns a 128-character value.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

A.8.5 cl_amd_compile_options

This extension adds the following options, which are not part of the OpenCL specification.

- `-g` — This is an experimental feature that lets you use the GNU project debugger, GDB, to debug kernels on x86 CPUs running Linux or cygwin/minGW under Windows. For more details, see Chapter 4, "Debugging OpenCL." This option does not affect the default optimization of the OpenCL code.
- `-O0` — Specifies to the compiler not to optimize. This is equivalent to the OpenCL standard option `-cl-opt-disable`.
- `-f[no-]bin-source` — Does [not] generate OpenCL source in the `.source` section. For more information, see Appendix C, "OpenCL Binary Image Format (BIF) v2.0."
- `-f[no-]bin-llvmir` — Does [not] generate LLVM IR in the `.llvmir` section. For more information, see Appendix C, "OpenCL Binary Image Format (BIF) v2.0."
- `-f[no-]bin-amdil` — Does [not] generate AMD IL in the `.amdil` section. For more information, see Appendix C, "OpenCL Binary Image Format (BIF) v2.0."
- `-f[no-]bin-exe` — Does [not] generate the executable (ISA) in .text section. For more information, see Appendix C, "OpenCL Binary Image Format (BIF) v2.0."

To avoid source changes, there are two environment variables that can be used to change CL options during the runtime.

- `AMD_OCL_BUILD_OPTIONS` — Overrides the CL options specified in `clBuildProgram()`.
- `AMD_OCL_BUILD_OPTIONS_APPEND` — Appends options to the options specified in `clBuildProgram()`.

A.8.6 cl_amd_offline_devices

To generate binary images offline, it is necessary to access the compiler for every device that the runtime supports, even if the device is currently not installed on the system. When, during context creation, `CL_CONTEXT_OFFLINE_DEVICES_AMD` is passed in the context properties, all supported devices, whether online or offline, are reported and can be used to create OpenCL binary images.

A.8.7 cl_amd_event_callback

This extension provides the ability to register event callbacks for states other than `cl_complete`. The full set of event states are allowed: `cl_queued`, `cl_submitted`, and `cl_running`. This extension is enabled automatically and does not need to be explicitly enabled through `#pragma` when using the SDK v2 of AMD Accelerated Parallel Processing.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

A.8.8 cl_amd_popcnt

This extension introduces a “population count” function called popcnt. This extension was taken into core OpenCL 1.2, and the function was renamed popcount. The core 1.2 popcount function (documented in section 6.12.3 of the *OpenCL Specification*) is identical to the AMD extension popcnt function.

A.8.9 cl_amd_media_ops

This extension adds the following built-in functions to the OpenCL language. Note: For OpenCL scalar types, *n* = 1; for vector types, it is {2, 4, 8, or 16}.

Note: in the following, ***n*** denotes the size, which can be 1, 2, 4, 8, or 16; [***i***] denotes the indexed element of a vector, designated 0 to ***n***-1.

Built-in function: amd_pack

```
uint amd_pack(float4 src)
```

Return value

```
((((uint)src[0]) & 0xFF) <<  0) +
((((uint)src[1]) & 0xFF) <<  8) +
((((uint)src[2]) & 0xFF) << 16) +
((((uint)src[3]) & 0xFF) << 24)
```

Built-in function: amd_unpack0

```
floatn   amd_unpack0 (uintn src)
```

Return value for each vector component

```
(float)(src[i] & 0xFF)
```

Built-in function: amd_unpack1

```
floatn   amd_unpack1 (uintn src)
```

Return value for each vector component

```
(float)((src[i] >> 8) & 0xFF)
```

Built-in function: amd_unpack2

```
floatn   amd_unpack2 (uintn src)
```

Return value for each vector component

```
(float)((src[i] >> 16) & 0xFF)
```

Built-in function: amd_unpack3

```
floatn  amd_unpack3(uintn src)
```

Return value for each vector component

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
(float)((src[i] >> 24) & 0xFF)
```

Built-in function: amd_bitalign

```
uintn  amd_bitalign (uintn src0, uintn src1, uintn src2)
```

Return value for each vector component

```
(uint) (((((long)src0[i]) << 32) | (long)src1[i]) >> (src2[i] & 31))
```

Built-in function: amd_bytealign

```
uintn  amd_bytealign (uintn src0, uintn src1, uintn src2)
```

Return value for each vector component

```
(uint) (((((long)src0[i]) << 32) | (long)src1[i]) >> ((src2[i] & 3)*8))
```

Built-in function: amd_lerp

```
uintn  amd_lerp (uintn src0, uintn src1, uintn src2)
```

Return value for each vector component

```
 (((((src0[i] >>  0) & 0xFF) + ((src1[i] >>  0) & 0xFF) + ((src2[i] >> 0) & 1)) >> 1) << 0) +
 (((((src0[i] >>  8) & 0xFF) + ((src1[i] >>  8) & 0xFF) + ((src2[i] >> 8) & 1)) >> 1) << 8) +
(((((src0[i] >> 16) & 0xFF) + ((src1[i] >> 16) & 0xFF) + ((src2[i] >> 16) & 1)) >> 1) << 16) +
(((((src0[i] >> 24) & 0xFF) + ((src1[i] >> 24) & 0xFF) + ((src2[i] >> 24) & 1)) >> 1) << 24) ;
```

Built-in function: amd_sad

```
uintn  amd_sad (uintn src0, uintn src1, uintn src2)
```

Return value for each vector component

```
src2[i] +
abs(((src0[i] >>  0) & 0xFF) - ((src1[i] >>  0) & 0xFF)) +
abs(((src0[i] >>  8) & 0xFF) - ((src1[i] >>  8) & 0xFF)) +
abs(((src0[i] >> 16) & 0xFF) - ((src1[i] >> 16) & 0xFF)) +
abs(((src0[i] >> 24) & 0xFF) - ((src1[i] >> 24) & 0xFF));
```

Built-in function: amd_sad4

```
uint amd_sad4 (uint4 a, uint4 b, uint c)
```

Return value for each vector component

```
src2[i] +
abs(((src0[i] >>  0) & 0xFF) - ((src1[i] >>  0) & 0xFF)) +
abs(((src0[i] >>  8) & 0xFF) - ((src1[i] >>  8) & 0xFF)) +
abs(((src0[i] >> 16) & 0xFF) - ((src1[i] >> 16) & 0xFF)) +
abs(((src0[i] >> 24) & 0xFF) - ((src1[i] >> 24) & 0xFF));
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Built-in function: amd_sadhi

```
uintn  amd_sadhi (uintn src0, uintn src1, uintn src2)
```

Return value for each vector component

```
src2[i] +
(abs(((src0[i] >>  0) & 0xFF) - ((src1[i] >>  0) & 0xFF)) << 16) +
(abs(((src0[i] >>  8) & 0xFF) - ((src1[i] >>  8) & 0xFF)) << 16) +
(abs(((src0[i] >> 16) & 0xFF) - ((src1[i] >> 16) & 0xFF)) << 16) +
(abs(((src0[i] >> 24) & 0xFF) - ((src1[i] >> 24) & 0xFF)) << 16);
```

For more information, see:
http://www.khronos.org/registry/cl/extensions/amd/cl_amd_media_ops.txt

A.8.10 cl_amd_media_ops2

This extension adds further built-in functions to those of cl_amd_media_ops. When enabled, it adds the following built-in functions to the OpenCL language.

Note: type*n* denotes an open scalar type { ***n*** = 1 } and vector types { ***n*** = 2, 4, 8, 16 }.

Built-in Function: uint*n* amd_msad (uint*n* src0, uint*n* src1, uint*n* src2)

Description:

```
uchar4 src0u8 = as_uchar4(src0.s0);
uchar4 src1u8 = as_uchar4(src1.s0);
dst.s0 = src2.s0 +
   ((src1u8.s0 == 0) ? 0 : abs(src0u8.s0 - src1u8.s0)) +
   ((src1u8.s1 == 0) ? 0 : abs(src0u8.s1 - src1u8.s1)) +
   ((src1u8.s2 == 0) ? 0 : abs(src0u8.s2 - src1u8.s2)) +
   ((src1u8.s3  == 0) ? 0  : abs(src0u8.s3  - src1u8.s3));
```

A similar operation is applied to other components of the vectors.

Built-in Function: ulong*n* amd_qsad (ulong*n* src0, uint*n* src1, ulong*n* src2)

Description:

```
uchar8 src0u8 = as_uchar8(src0.s0);
ushort4 src2u16 = as_ushort4(src2.s0);
ushort4 dstu16;
dstu16.s0 = amd_sad(as_uint(src0u8.s0123), src1.s0, src2u16.s0);
dstu16.s1 = amd_sad(as_uint(src0u8.s1234), src1.s0, src2u16.s1);
dstu16.s2 = amd_sad(as_uint(src0u8.s2345), src1.s0, src2u16.s2);
dstu16.s3 = amd_sad(as_uint(src0u8.s3456), src1.s0, src2u16.s3);
dst.s0     = as_uint2(dstu16);
```

A similar operation is applied to other components of the vectors.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Built-in Function:

```
ulongn amd_mqsad (ulongn src0, uintn src1, ulongn src2)
```

Description:

```
uchar8 src0u8 = as_uchar8(src0.s0);
ushort4 src2u16 = as_ushort4(src2.s0);
ushort4 dstu16;
dstu16.s0 = amd_msad(as_uint(src0u8.s0123), src1.s0, src2u16.s0);
dstu16.s1 = amd_msad(as_uint(src0u8.s1234), src1.s0, src2u16.s1);
dstu16.s2 = amd_msad(as_uint(src0u8.s2345), src1.s0, src2u16.s2);
dstu16.s3 = amd_msad(as_uint(src0u8.s3456), src1.s0, src2u16.s3);
dst.s0     = as_uint2(dstu16);
```

A similar operation is applied to other components of the vectors.

Built-in Function: `uintn amd_sadw (uintn src0, uintn src1, uintn src2)`

Description:

```
ushort2 src0u16 = as_ushort2(src0.s0);
ushort2 src1u16 = as_ushort2(src1.s0);
dst.s0 = src2.s0 +
   abs(src0u16.s0 - src1u16.s0) +
   abs(src0u16.s1 - src1u16.s1);
```

A similar operation is applied to other components of the vectors.

Built-in Function: `uintn amd_sadd (uintn src0, uintn src1, uintn src2)`

Description:

```
dst.s0 = src2.s0 +  abs(src0.s0 - src1.s0);
```

A similar operation is applied to other components of the vectors.

Built-in Function: `uintn amd_bfm (uintn src0, uintn src1)`

Description:

```
dst.s0 = ((1 << (src0.s0 & 0x1f)) - 1) << (src1.s0 & 0x1f);
```

A similar operation is applied to other components of the vectors.

Built-in Function: `uintn amd_bfe (uintn src0, uintn src1, uintn src2)`

Description:

NOTE: The **>>** operator represents a logical right shift.

```
offset = src1.s0 & 31;
width = src2.s0 & 31;
if width = 0
   dst.s0 = 0;
else if (offset + width) < 32
   dst.s0 = (src0.s0 << (32 - offset - width)) >> (32 - width);
else
   dst.s0 = src0.s0 >> offset;
```

A similar operation is applied to other components of the vectors.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Built-in Function: `intn amd_bfe (intn src0, uintn src1, uintn src2)`

Description:

NOTE: operator **>>** represent an arithmetic right shift.

```
offset = src1.s0 & 31;
width = src2.s0 & 31;
if width = 0
   dst.s0 = 0;
else if (offset + width) < 32
   dst.s0 = src0.s0 << (32-offset-width) >> 32-width;
else
   dst.s0 = src0.s0 >> offset;
```

A similar operation is applied to other components of the vectors.

Built-in Function:

```
intn amd_median3 (intn src0, intn src1, intn src2)
uintn amd_median3 (uintn src0, uintn src1, uintn src2)
floatn amd_median3 (floatn src0, floatn src1, floatn src2)
```

Description:

Returns median of src0, src1, and src2.

Built-in Function:

```
intn amd_min3 (intn src0, intn src1, intn src2)
uintn amd_min3 (uintn src0, uintn src1, uintn src2)
floatn amd_min3 (floatn src0, floatn src1, floatn src2)
```

Description:

Returns min of src0, src1, and src2.

Built-in Function:

```
intn amd_max3 (intn src0, intn src1, intn src2)
uintn amd_max3 (uintn src0, uintn src1, uintn src2)
floatn amd_max3 (floatn src0, floatn src1, floatn src2)
```

Description:

Returns max of src0, src1, and src2.

For more information, see:
http://www.khronos.org/registry/cl/extensions/amd/cl_amd_media_ops2.txt

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

A.8.11 cl_amd_printf

The OpenCL™ Specification 1.1 and 1.2 support the optional AMD extension `cl_amd_printf`, which provides printf capabilities to OpenCL C programs. To use this extension, an application first must include `#pragma OPENCL EXTENSION cl_amd_printf : enable`.

Built-in function:

```
printf(__constant char * restrict format, …);
```

This function writes output to the `stdout` stream associated with the host application. The `format` string is a character sequence that:

– is null-terminated and composed of zero and more directives,

– ordinary characters (i.e. not %), which are copied directly to the output stream unchanged, and

– conversion specifications, each of which can result in fetching zero or more arguments, converting them, and then writing the final result to the output stream.

The `format` string must be resolvable at compile time; thus, it cannot be dynamically created by the executing program. (Note that the use of variadic arguments in the built-in `printf` does not imply its use in other built-ins; more importantly, it is not valid to use `printf` in user-defined functions or kernels.)

The OpenCL C printf closely matches the definition found as part of the C99 standard. Note that conversions introduced in the `format` string with `%` are supported with the following guidelines:

- A 32-bit floating point argument is not converted to a 64-bit double, unless the extension `cl_khr_fp64` is supported and enabled, as defined in section 9.3 of the *OpenCL Specification 1.1*. This includes the double variants if `cl_khr_fp64` is supported and defined in the corresponding compilation unit.
- 64-bit integer types can be printed using `%ld / %lx / %lu` .
- `%lld / %llx / %llu` are not supported and reserved for 128-bit integer types (long long).
- All OpenCL vector types (section 6.1.2 of the *OpenCL Specification 1.1*) can be explicitly passed and printed using the modifier `v`*`n`*, where *`n`* can be 2, 3, 4, 8, or 16. This modifier appears before the original conversion specifier for the vector's component type (for example, to print a `float4 %v4f`). Since `v`*`n`* is a conversion specifier, it is valid to apply optional flags, such as `field width` and `precision`, just as it is when printing the component types. Since a vector is an aggregate type, the comma separator is used between the components: `0:1, … , n-2:n-1`.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

A.8.12 cl_amd_predefined_macros

The following macros are predefined when compiling OpenCL™ C kernels. These macros are defined automatically based on the device for which the code is being compiled.

GPU devices:

```
__WinterPark__
__BeaverCreek__
__Turks__
__Caicos__
__Tahiti__
__Pitcairn__
__Capeverde__
__Cayman__
__Barts__
__Cypress__
__Juniper__
__Redwood__
__Cedar__
__ATI_RV770__
__ATI_RV730__
__ATI_RV710__
__Loveland__
__GPU__
```

CPU devices:

```
__CPU__
__X86__
__X86_64__
```

Note that `__GPU__` or `__CPU__` are predefined whenever a GPU or CPU device is the compilation target.

An example kernel is provided below.

```
#pragma OPENCL EXTENSION cl_amd_printf : enable
const char* getDeviceName() {
#ifdef __Cayman__
        return "Cayman";
#elif __Barts__
        return "Barts";
#elif __Cypress__
        return "Cypress";
#elif defined(__Juniper__)
        return "Juniper";
#elif defined(__Redwood__)
        return "Redwood";
#elif defined(__Cedar__)
        return "Cedar";
#elif defined(__ATI_RV770__)
        return "RV770";
#elif defined(__ATI_RV730__)
        return "RV730";
#elif defined(__ATI_RV710__)
        return "RV710";
#elif defined(__Loveland__)
        return "Loveland";
#elif defined(__GPU__)
        return "GenericGPU";
#elif defined(__X86__)
        return "X86CPU";
#elif defined(__X86_64__)
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
        return "X86-64CPU";
#elif defined(__CPU__)
        return "GenericCPU";
#else
        return "UnknownDevice";
#endif
}
kernel void test_pf(global int* a)
{
        printf("Device Name: %s\n", getDeviceName());
}
```

A.8.13 cl_amd_bus_addressable_memory

This extension defines an API for peer-to-peer transfers between AMD GPUs and other PCIe device, such as third-party SDI I/O devices. Peer-to-peer transfers have extremely low latencies by not having to use the host's main memory or the CPU (see Figure A.1). This extension allows sharing a memory allocated by the graphics driver to be used by other devices on the PCIe bus (peer-to-peer transfers) by exposing a write-only bus address. It also allows memory allocated on other PCIe devices (non-AMD GPU) to be directly accessed by AMD GPUs. One possible use of this is for a video capture device to directly write into the GPU memory using its DMA.This extension is supported only on AMD FirePro™ professional graphics cards.

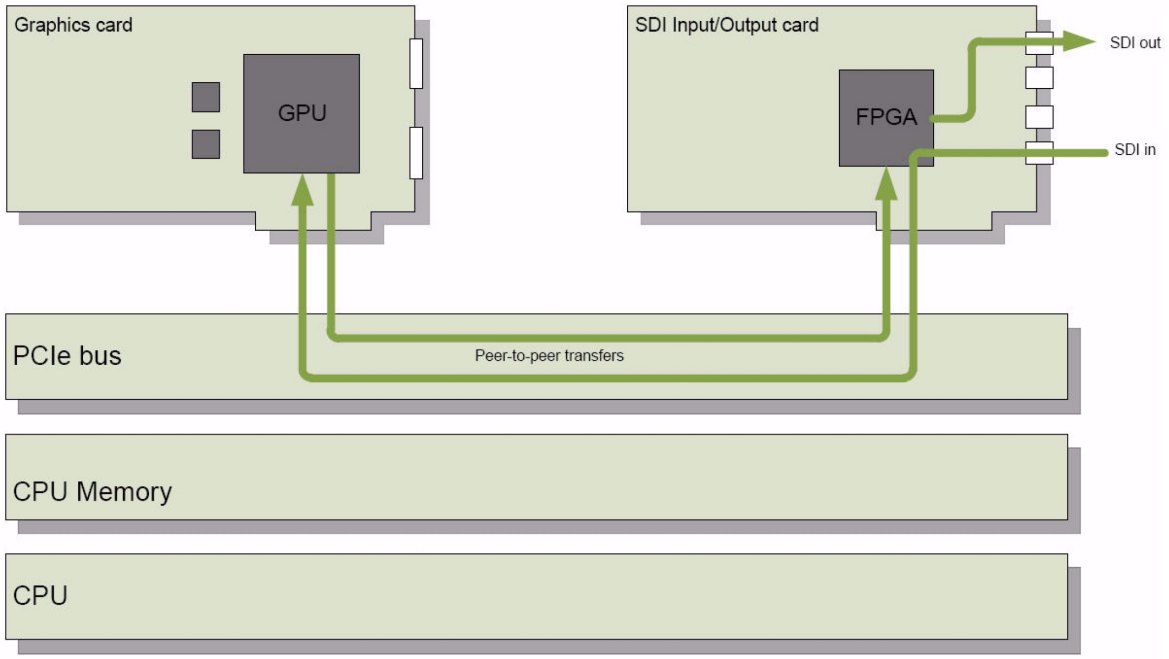

Figure A.1 Peer-to-Peer Transfers Using the cl_amd_bus_addressable_memory Extension

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

A.9 Supported Functions for cl_amd_fp64 / cl_khr_fp64

AMD OpenCL is now cl_khr_fp64-compliant on devices compliant with OpenCL 1.1 and greater. Thus, cl_amd_fp64 is now a synonym for cl_khr_fp64 on all supported devices.

A.10 Extension Support by Device

Table A.1 and Table A.2 list the extension support for selected devices.

Table A.1 Extension Support for AMD GPU Devices 1

	AMD APUs			AMD Radeon™ HD				
Extension	Brazos	Llano	Trinity	Tahiti[1], Pitcairn[2], Cape Verde[3]	Turks[4]	Cayman[5]	Barts[6]	Cypress[7]
cl_khr_*_atomics (32-bit)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
cl_ext_atomic_counters_32	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
cl_khr_gl_sharing	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
cl_khr_byte_addressable_store	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
cl_ext_device_fission	CPU only	CPU only	CPU only	No	No	No	No	No
cl_amd_device_attribute_query	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
cl_khr_fp64	CPU only	CPU only	CPU only	Yes	Yes	Yes	No	Yes
cl_amd_fp64	CPU only	CPU only	CPU only	Yes	Yes	Yes	No	Yes
cl_amd_vec3	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Images	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
cl_khr_d3d10_sharing	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
cl_amd_media_ops	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
cl_amd_printf	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
cl_amd_popcnt	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
cl_khr_3d_image_writes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Platform Extensions								
cl_khr_icd	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
cl_amd_event_callback	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
cl_amd_offline_devices	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

1. AMD Radeon™ HD 79XX series.
2. AMD Radeon™ HD 78XX series.
3. AMD Radeon™ HD 77XX series.
4. AMD Radeon™ HD 75XX series and AMD Radeon™ HD 76XX series.
5. AMD Radeon™ HD 69XX series.
6. AMD Radeon™ HD 68XX series.
7. ATI Radeon™ HD 59XX series and 58XX series, AMD FirePro™ V88XX series and V87XX series.

Note that an atomic counter is a device-level counter that can be added / decremented by different work-items, where the atomicity of the operation is

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

guaranteed. The access to the counter is done only through add/dec built-in functions; thus, no two work-items have the same value returned in the case that a given kernel only increments or decrements the counter. (Also see http://www.khronos.org/registry/cl/extensions/ext/cl_ext_atomic_counters_32.txt.)

Table A.2 Extension Support for Older AMD GPUs and CPUs

Extension	Juniper[1]	Redwood[2]	Cedar[3]	x86 CPU with SSE2 or later
cl_khr_*_atomics	Yes	Yes	Yes	Yes
cl_ext_atomic_counters_32	Yes	Yes	Yes	No
cl_khr_gl_sharing	Yes	Yes	Yes	Yes
cl_khr_byte_addressable_store	Yes	Yes	Yes	Yes
cl_ext_device_fission	No	No	No	Yes
cl_amd_device_attribute_query	Yes	Yes	Yes	Yes
cl_khr_fp64	No	No	No	Yes
cl_amd_fp64[4]	No	No	No	Yes
cl_amd_vec3	Yes	Yes	Yes	Yes
Images	Yes	Yes	Yes	Yes[5]
cl_khr_d3d10_sharing	Yes	Yes	Yes	Yes
cl_amd_media_ops	Yes	Yes	Yes	Yes
cl_amd_media_ops2	Yes	Yes	Yes	Yes
cl_amd_printf	Yes	Yes	Yes	Yes
cl_amd_popcnt	Yes	Yes	Yes	Yes
cl_khr_3d_image_writes	Yes	Yes	Yes	No
Platform Extensions				
cl_khr_icd	Yes	Yes	Yes	Yes
cl_amd_event_callback	Yes	Yes	Yes	Yes
cl_amd_offline_devices	Yes	Yes	Yes	Yes

1. ATI Radeon™ HD 5700 series, AMD Mobility Radeon™ HD 5800 series, AMD FirePro™ V5800 series, AMD Mobility FirePro™ M7820.
2. ATI Radeon™ HD 5600 Series, ATI Radeon™ HD 5600 Series, ATI Radeon™ HD 5500 Series, AMD Mobility Radeon™ HD 5700 Series, AMD Mobility Radeon™ HD 5600 Series, AMD FirePro™ V4800 Series, AMD FirePro™ V3800 Series, AMD Mobility FirePro™ M5800
3. ATI Radeon™ HD 5400 Series, AMD Mobility Radeon™ HD 5400 Series
4. Available on all devices that have double-precision, including all Southern Island devices.
5. Environment variable `CPU_IMAGE_SUPPORT` must be set.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Appendix B
The OpenCL Installable Client Driver (ICD)

The OpenCL Installable Client Driver (ICD) is part of the AMD Accelerated Parallel Processing software stack. Code written prior to SDK v2.0 must be changed to comply with OpenCL ICD requirements.

B.1 Overview

The ICD allows multiple OpenCL implementations to co-exist; also, it allows applications to select between these implementations at runtime.

Use the `clGetPlatformIDs()` and `clGetPlatformInfo()` functions to see the list of available OpenCL implementations, and select the one that is best for your requirements. It is recommended that developers offer their users a choice on first run of the program or whenever the list of available platforms changes.

A properly implemented ICD and OpenCL library is transparent to the end-user.

B.2 Using ICD

Sample code that is part of the SDK contains examples showing how to query the platform API and call the functions that require a valid platform parameter.

This is a pre-ICD code snippet.

```
context = clCreateContextFromType(
            0,
            dType,
            NULL,
            NULL,
            &status);
```

The ICD-compliant version of this code follows.

```
/*
     * Have a look at the available platforms and pick either
     * the AMD one if available or a reasonable default.
     */


    cl_uint numPlatforms;
    cl_platform_id platform = NULL;
    status = clGetPlatformIDs(0, NULL, &numPlatforms);
    if(!sampleCommon->checkVal(status,
                               CL_SUCCESS,
                               "clGetPlatformIDs failed."))
    {
        return SDK_FAILURE;
    }
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
    if (0 < numPlatforms)
    {
        cl_platform_id* platforms = new cl_platform_id[numPlatforms];
        status = clGetPlatformIDs(numPlatforms, platforms, NULL);
        if(!sampleCommon->checkVal(status,
                                   CL_SUCCESS,
                                   "clGetPlatformIDs failed."))
        {
            return SDK_FAILURE;
        }
        for (unsigned i = 0; i < numPlatforms; ++i)
        {
            char pbuf[100];
            status = clGetPlatformInfo(platforms[i],
                                       CL_PLATFORM_VENDOR,
                                       sizeof(pbuf),
                                       pbuf,
                                       NULL);


            if(!sampleCommon->checkVal(status,
                                       CL_SUCCESS,
                                       "clGetPlatformInfo failed."))
            {
                return SDK_FAILURE;
            }


            platform = platforms[i];
            if (!strcmp(pbuf, "Advanced Micro Devices, Inc."))
            {
                break;
            }
        }
        delete[] platforms;
    }


    /*
     * If we could find our platform, use it. Otherwise pass a NULL and
get whatever the
     * implementation thinks we should be using.
     */


    cl_context_properties cps[3] =
    {
        CL_CONTEXT_PLATFORM,
        (cl_context_properties)platform,
        0
    };
    /* Use NULL for backward compatibility */
    cl_context_properties* cprops = (NULL == platform) ? NULL : cps;


    context = clCreateContextFromType(
                  cprops,
                  dType,
                  NULL,
                  NULL,
                  &status);
```

Another example of a pre-ICD code snippet follows.

```
status = clGetDeviceIDs(NULL, CL_DEVICE_TYPE_DEFAULT, 0, NULL,
&numDevices);
```

The ICD-compliant version of the code snippet is:

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
status = clGetDeviceIDs(platform, CL_DEVICE_TYPE_DEFAULT, 0, NULL,
&numDevices);
```

NOTE: It is recommended that the host code look at the platform vendor string when searching for the desired OpenCL platform, instead of using the platform name string. The platform name string might change, whereas the platform vendor string remains constant for a particular vendor's implementation.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Appendix C
OpenCL Binary Image Format (BIF) v2.0

C.1 Overview

OpenCL Binary Image Format (BIF) 2.0 is in the ELF format. BIF2.0 allows the OpenCL binary to contain the OpenCL source program, the LLVM IR, and the executable. The BIF defines the following special sections:

- `.source`: for storing the OpenCL source program.
- `.llvmir`: for storing the OpenCL immediate representation (LLVM IR).
- `.comment`: for storing the OpenCL version and the driver version that created the binary.

The BIF can have other special sections for debugging, etc. It also contains several ELF special sections, such as:

- `.text` for storing the executable.
- `.rodata` for storing the OpenCL runtime control data.
- other ELF special sections required for forming an ELF (for example: `.strtab`, `.symtab`, `.shstrtab`).

By default, OpenCL generates a binary that has LLVM IR, and the executable for the GPU (,`.llvmir`, `.amdil`, and `.text` sections), as well as LLVM IR and the executable for the CPU (`.llvmir` and `.text` sections). The BIF binary always contains a `.comment` section, which is a readable C string. The default behavior can be changed with the BIF options described in Section C.2, “BIF Options,” page C-3.

The LLVM IR enables recompilation from LLVM IR to the target. When a binary is used to run on a device for which the original program was not generated and the original device is feature-compatible with the current device, OpenCL recompiles the LLVM IR to generate a new code for the device. Note that the LLVM IR is only universal within devices that are feature-compatible in the same device type, not across different device types. This means that the LLVM IR for the CPU is not compatible with the LLVM IR for the GPU. The LLVM IR for a GPU works only for GPU devices that have equivalent feature sets.

BIF2.0 is supported since Stream SDK 2.2.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

C.1.1 Executable and Linkable Format (ELF) Header

For the ELF binary to be considered valid, the AMD OpenCL runtime expects certain values to be specified. The following header fields must be set for all binaries that are created outside of the OpenCL framework.

Table C.1 ELF Header Fields

Field	Value	Description
e_ident[EI_CLASS]	ELFCLASS32, ELFCLASS64	BIF can be either 32-bit ELF or 64bit ELF.
e_ident[EI_DATA]	ELFDATA2LSB	BIF is stored in little Endian order.
e_ident[EI_OSABI]	ELFOSABI_NONE	Not used.
e_ident[EI_ABIVERSION]	0	Not used.
e_type	ET_NONE	Not used.
e_machine	oclElfTargets Enum	CPU/GPU machine ID.
E_version	EV_CURRENT	Must be EV_CURRENT.
e_entry	0	Not used.
E_phoff	0	Not used.
e_flags	0	Not used.
E_phentsize	0	Not used.
E_phnum	0	Not used.

The fields not shown in Table C.1 are given values according to the *ELF Specification*. The e_machine value is defined as one of the oclElfTargets enumerants; the values for these are:

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

e_machine = { 1001 + CaltargetEnum : GPU; 2002 : CPU generic without SSE3; 2003 : CPU generic with SSE3

```
typedef enum CALtargetEnum {
    CAL_TARGET_600 = 0, /**< R600 GPU ISA */
    CAL_TARGET_610 = 1, /**< RV610 GPU ISA */
    CAL_TARGET_630 = 2, /**< RV630 GPU ISA */
    CAL_TARGET_670 = 3, /**< RV670 GPU ISA */
    CAL_TARGET_7XX = 4, /**< R700 class GPU ISA */
    CAL_TARGET_770 = 5, /**< RV770 GPU ISA */
    CAL_TARGET_710 = 6, /**< RV710 GPU ISA */
    CAL_TARGET_730 = 7, /**< RV730 GPU ISA */
    CAL_TARGET_CYPRESS = 8, /**< CYPRESS GPU ISA */
    CAL_TARGET_JUNIPER = 9, /**< JUNIPER GPU ISA */
    CAL_TARGET_REDWOOD = 10, /**< REDWOOD GPU ISA */
    CAL_TARGET_CEDAR= 11, /**< CEDAR GPU ISA */
    CAL_TARGET_SUMO = 12, /**< SUMO GPU ISA */
    CAL_TARGET_SUPERSUMO =13, /**< SUPERSUMO GPU ISA */
    CAL_TARGET_WRESTLER = 14, /**< WRESTLER GPU ISA */
    CAL_TARGET_CAYMAN =15, /**< CAYMAN GPU ISA */
    CAL_TARGET_KAUAI = 16, /**< KAUAI GPU ISA */
    CAL_TARGET_BARTS = 17 , /**< BARTS GPU ISA */
    CAL_TARGET_TURKS = 18 , /**< TURKS GPU ISA */
    CAL_TARGET_CAICOS  = 19, /**< CAICOS GPU ISA */
    CAL_TARGET_TAHITI = 20,/**< TAHITI GPU ISA*/
    CAL_TARGET_PITCAIRN = 21,/**< PITCAIRN GPU ISA*/
    CAL_TARGET_CAPEVERDE = 22,/**< CAPE VERDE GPU ISA*/
    CAL_TARGET_DEVASTATOR = 23,/**< DEVASTATOR GPU ISA*/
    CAL_TARGET_SCRAPPER = 24, /**< SCRAPPER GPU ISA*/
    CAL_TARGET_OLAND = 25, /**< OLAND GPU ISA*/
    CAL_TARGET_BONAIRE = 26, /**< BONAIRE GPU ISA*/
    CAL_TARGET_KALINDI = 29, /**< KALINDI GPU ISA*/
};
```

C.1.2 Bitness

The BIF can be either 32-bit ELF format or a 64-bit ELF format. For the GPU, OpenCL generates a 32-bit BIF binary; it can read either 32-bit BIF or 64-bit BIF binary. For the CPU, OpenCL generates and reads only 32-bit BIF binaries if the host application is 32-bit (on either 32-bit OS or 64-bit OS). It generates and reads only 64-bit BIF binary if the host application is 64-bit (on 64-bit OS).

C.2 BIF Options

OpenCL provides the following options to control what is contained in the binary.

`-f[no-]bin-source` — [not] generate OpenCL source in `.source` section.

`-f[no-]bin-llvmir` — [not] generate LLVM IR in `.llvmir` section.

`-f[no-]bin-exe` — [not] generate the executable (ISA) in `.text` section.

The option syntax follows the GCC option syntax.

By default, OpenCL generates the `.llvmir` section, `.amdil` section, and `.text` section. The following are examples for using these options:

Example 1: Generate executable for execution:

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
clBuildProgram(program, 0, NULL, "-fno-bin-llvmir -fno-bin-amdil", NULL,
NULL);
```

Example 2: Generate only LLVM IR:

```
clBuildProgram(program, 0, NULL, "-fno-bin-exe -fno-bin-amdil", NULL,
NULL);
```

This binary can recompile for all the other devices of the same device type.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Appendix D
Hardware overview of pre-GCN devices

This chapter provides a hardware overview of pre-GCN devices. Pre-GCN devices include the Evergreen and Northern Islands families that are based on VLIW.

A general OpenCL device comprises compute units, each of which can have multiple processing elements. A work-item (or SPMD kernel instance) executes on a single processing element. The processing elements within a compute unit can execute in lock-step using SIMD execution. Compute units, however, execute independently (see Figure D.1).

AMD GPUs consist of multiple compute units. The number of them and the way they are structured varies with the device family, as well as device designations within a family. Each of these processing elements possesses ALUs. For devices in the Northern Islands and Southern Islands families, these ALUs are arranged in four (in the Evergreen family, there are five) processing elements with arrays of 16 ALUs. Each of these arrays executes a single instruction across each lane for each of a block of 16 work-items. That instruction is repeated over four cycles to make the 64-element vector called a *wavefron*t. On Northern Islands and Evergreen family devices, the PE arrays execute instructions from one wavefront, so that each work-item issues four (for Northern Islands) or five (for Evergreen) instructions at once in a very-long-instruction-word (VLIW) packet.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Figure D.1 shows a simplified block diagram of a generalized AMD GPU compute device.

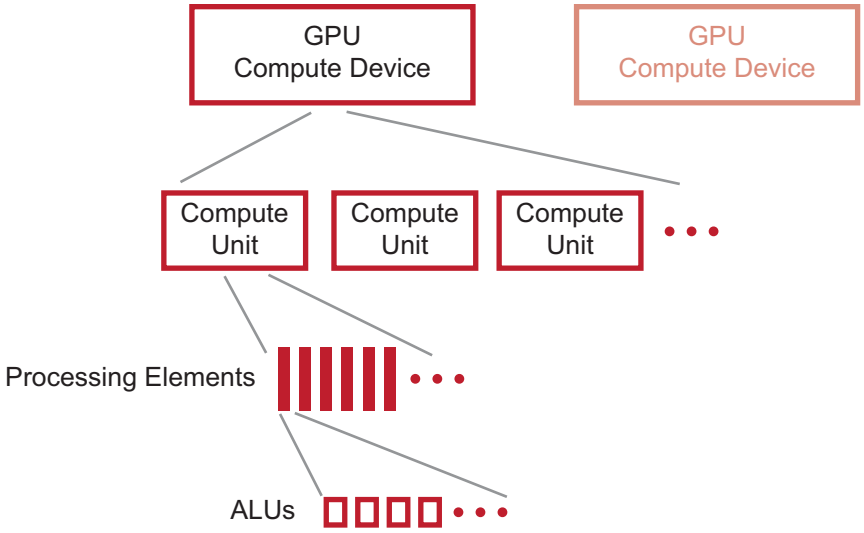

Figure D.1 Generalized AMD GPU Compute Device Structure

Figure D.2 is a simplified diagram of an AMD GPU compute device. Different GPU compute devices have different characteristics (such as the number of compute units), but follow a similar design pattern.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

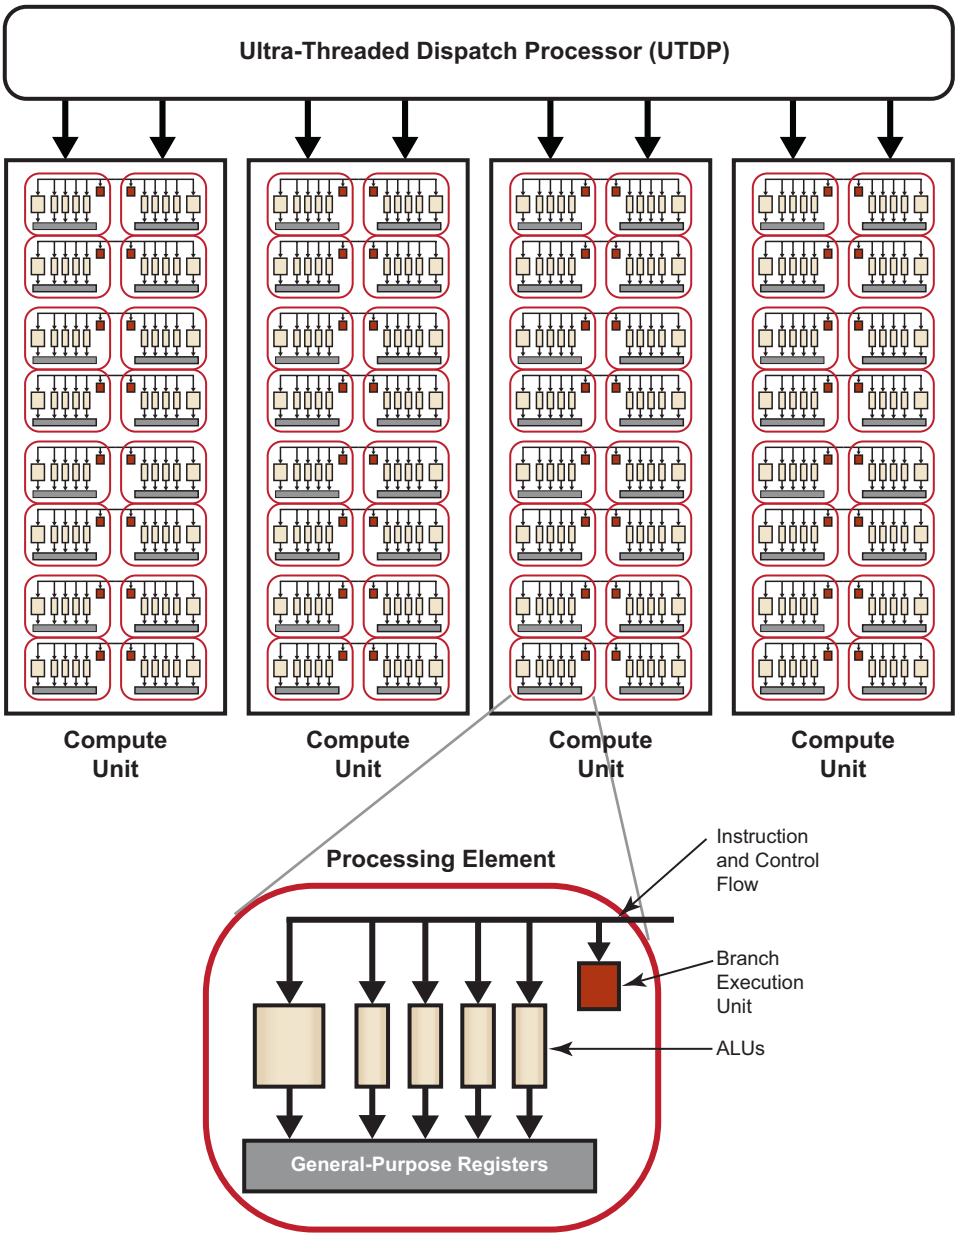

Figure D.2 Simplified Block Diagram of an Evergreen-Family GPU[1]

GPU compute devices comprise groups of compute units. Each compute unit contains numerous processing elements, which are responsible for executing kernels, each operating on an independent data stream. Processing elements, in turn, contain numerous processing elements, which are the fundamental, programmable ALUs that perform integer, single-precision floating-point, double-precision floating-point, and transcendental operations. All processing elements within a compute unit execute the same instruction sequence in lock-step for Evergreen and Northern Islands devices; different compute units can execute

1. Much of this is transparent to the programmer.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

different instructions.

A processing element is arranged as a five-way or four-way (depending on the GPU type) very long instruction word (VLIW) processor (see bottom of Figure D.2). Up to five scalar operations (or four, depending on the GPU type) can be co-issued in a VLIW instruction, each of which are executed on one of the corresponding five ALUs. ALUs can execute single-precision floating point or integer operations. One of the five ALUs also can perform transcendental operations (sine, cosine, logarithm, etc.). Double-precision floating point operations are processed (where supported) by connecting two or four of the ALUs (excluding the transcendental core) to perform a single double-precision operation. The processing element also contains one branch execution unit to handle branch instructions.

Different GPU compute devices have different numbers of processing elements. For example, the ATI Radeon™ HD 5870 GPU has 20 compute units, each with 16 processing elements, and each processing elements contains five ALUs; this yields 1600 physical ALUs.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Appendix E
OpenCL-OpenGL Interoperability

This chapter explains how to establish an association between GL context and CL context.

Please note the following guidelines.

1. All devices used to create the OpenCL context associated with `command_queue` must support acquiring shared CL/GL objects. This constraint is enforced at context-creation time.
2. `clCreateContext` and `clCreateContextFromType` fail context creation if the device list passed in cannot interoperate with the GL context. `clCreateContext` only permits GL-friendly device(s). `clCreateFromContextType` can only include GL-friendly device(s).
3. Use `clGetGLContextInfoKHR` to determine GL-friendly device(s) from the following parameters:
   a. `CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR` only returns the device that can interoperate with the GL context.
   b. `CL_DEVICES_FOR_GL_CONTEXT_KHR` includes all GL-context interoperable devices.
4. While it is possible to create as many GL contexts on a GPU, do not create concurrently two GL contexts for two GPUs from the same process.
5. For OpenGL interoperability with OpenCL, there is a strict order in which the OpenCL context is created and the texture/buffer shared allocations can be made. To use shared resources, the OpenGL application must create an OpenGL context and afterwards an OpenCL context. All resources (GL buffers and textures) created after the OpenCL context was created can be shared between OpenGL and OpenCL. If resources are allocated before the OpenCL context was created, they cannot be shared between OpenGL and OpenCL.

E.1 Under Windows

This sections discusses CL-GL interoperability for single and multiple GPU systems running under Windows.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

E.1.1 Single GPU Environment

5.1.1.1 Creating CL Context from a GL Context

Use **GLUT** windowing system or Win32 API for event handling.

Using GLUT

1. Use `glutInit` to initialize the GLUT library and negotiate a session with the windowing system. This function also processes the command line options, depending on the windowing system.
2. Use `wglGetCurrentContext` to get the current rendering GL context (HGLRC) of the calling thread.
3. Use `wglGetCurrentDC` to get the device context (HDC) that is associated with the current OpenGL rendering context of the calling thread.
4. Use the `clGetGLContextInfoKHR` *(*See Section 9.7 of the *OpenCL Specification 1.1)* function and the `CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR` parameter to get the device ID of the CL device associated with OpenGL context.
5. Use `clCreateContext` *(*See Section 4.3 of the *OpenCL Specification 1.1)* to create the CL context (of type `cl_context`).

The following code snippet shows you how to create an interoperability context using GLUT on single GPU system.

```
glutInit(&argc, argv);
glutInitDisplayMode(GLUT_RGBA | GLUT_DOUBLE);
glutInitWindowSize(WINDOW_WIDTH, WINDOW_HEIGHT);
glutCreateWindow("OpenCL SimpleGL");

HGLRC glCtx = wglGetCurrentContext();

Cl_context_properties cpsGL[] =
{CL_CONTEXT_PLATFORM,(cl_context_properties)platform,
CL_WGL_HDC_KHR, (intptr_t) wglGetCurrentDC(),

CL_GL_CONTEXT_KHR, (intptr_t) glCtx, 0};

status = clGetGLContextInfoKHR(cpsGL,
               CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR,
               sizeof(cl_device_id),
               &interopDevice,
               NULL);

// Create OpenCL context from device's id
               context = clCreateContext(cpsGL,
               1,
               &interopDevice,
               0,
               0,
               &status);
```

Using Win32 API

1. Use `CreateWindow` for window creation and get the device handle (HWND).

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

2. Use `GetDC` to get a handle to the device context for the client area of a specific window, or for the entire screen (OR). Use `CreateDC` function to create a device context (HDC) for the specified device.
3. Use `ChoosePixelFormat` to match an appropriate pixel format supported by a device context and to a given pixel format specification.
4. Use `SetPixelFormat` to set the pixel format of the specified device context to the format specified.
5. Use `wglCreateContext` to create a new OpenGL rendering context from device context (HDC).
6. Use `wglMakeCurrent` to bind the GL context created in the above step as the current rendering context.
7. Use `clGetGLContextInfoKHR` function (see Section 9.7 of the *OpenCL Specification 1.1*) and parameter `CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR` to get the device ID of the CL device associated with OpenGL context.
8. Use `clCreateContext` function (see Section 4.3 of the *OpenCL Specification 1.1*) to create the CL context (of type `cl_context`).

The following code snippet shows how to create an interoperability context using WIN32 API for windowing. (Users also can refer to the SimpleGL sample in the AMD APP SDK samples.)

```
int  pfmt;
PIXELFORMATDESCRIPTOR  pfd;
pfd.nSize           = sizeof(PIXELFORMATDESCRIPTOR);
pfd.nVersion        = 1;
pfd.dwFlags         = PFD_DRAW_TO_WINDOW |
                          PFD_SUPPORT_OPENGL | PFD_DOUBLEBUFFER ;
pfd.iPixelType      = PFD_TYPE_RGBA;
pfd.cColorBits      = 24;
pfd.cRedBits        = 8;
pfd.cRedShift       = 0;
pfd.cGreenBits      = 8;
pfd.cGreenShift     = 0;
pfd.cBlueBits       = 8;
pfd.cBlueShift      = 0;
pfd.cAlphaBits      = 8;
pfd.cAlphaShift     = 0;
pfd.cAccumBits      = 0;
pfd.cAccumRedBits   = 0;
pfd.cAccumGreenBits = 0;
pfd.cAccumBlueBits  = 0;
pfd.cAccumAlphaBits = 0;
pfd.cDepthBits      = 24;
pfd.cStencilBits    = 8;
pfd.cAuxBuffers     = 0;
pfd.iLayerType      = PFD_MAIN_PLANE;
pfd.bReserved       = 0;
pfd.dwLayerMask     = 0;
pfd.dwVisibleMask   = 0;
pfd.dwDamageMask    = 0;

ZeroMemory(&pfd, sizeof(PIXELFORMATDESCRIPTOR));

WNDCLASS windowclass;

windowclass.style = CS_OWNDC;
windowclass.lpfnWndProc = WndProc;
windowclass.cbClsExtra = 0;
windowclass.cbWndExtra = 0;
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
windowclass.hInstance = NULL;
windowclass.hIcon = LoadIcon(NULL, IDI_APPLICATION);
windowclass.hCursor = LoadCursor(NULL, IDC_ARROW);
windowclass.hbrBackground = (HBRUSH)GetStockObject(BLACK_BRUSH);
windowclass.lpszMenuName = NULL;
windowclass.lpszClassName = reinterpret_cast<LPCSTR>("SimpleGL");
         RegisterClass(&windowclass);

gHwnd = CreateWindow(reinterpret_cast<LPCSTR>("SimpleGL"),
           reinterpret_cast<LPCSTR>("SimpleGL"),
           WS_CAPTION | WS_POPUPWINDOW | WS_VISIBLE,
           0,
           0,
           screenWidth,
           screenHeight,
           NULL,
           NULL,
           windowclass.hInstance,
           NULL);

hDC = GetDC(gHwnd);

pfmt = ChoosePixelFormat(hDC, &pfd);

ret = SetPixelFormat(hDC, pfmt, &pfd);

hRC = wglCreateContext(hDC);

ret = wglMakeCurrent(hDC, hRC);

cl_context_properties properties[] =
{
    CL_CONTEXT_PLATFORM,
    (cl_context_properties) platform,
    CL_GL_CONTEXT_KHR,   (cl_context_properties) hRC,
    CL_WGL_HDC_KHR,      (cl_context_properties) hDC,
    0
};

status = clGetGLContextInfoKHR(properties,
           CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR,
           sizeof(cl_device_id),
           &interopDevice,
           NULL);

// Create OpenCL context from device's id
context = clCreateContext(properties,
           1,
           &interopDevice,
           0,
           0,
           &status);
```

E.1.2 Multi-GPU Environment

5.1.2.1 Creating CL context from a GL context

Do not to use the GLUT windowing system in multi-GPU environment because it always creates a GL context on the primary display, and it is not possible to specify which display device to select for a GL context.

To use Win32 API for windowing in multi-GPU environment:

1. Detect each display by using `EnumDisplayDevices` function. This function lets you obtain the information about display devices in the current session.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

2. To query all display devices in the current session, call this function in a loop, starting with DevNum set to 0, and incrementing DevNum until the function fails. To select all display devices in the desktop, use only the display devices that have the DISPLAY_DEVICE_ATTACHED_TO_DESKTOP flag in the DISPLAY_DEVICE structure.
3. To get information on the display adapter, call EnumDisplayDevices with lpDevice set to NULL. For example, DISPLAY_DEVICE.**DeviceString** contains the adapter name.
4. Use EnumDisplaySettings to get DEVMODE. dmPosition.x and dmPosition.y are used to get the x coordinate and y coordinate of the current display.
5. Try to find the first OpenCL device (winner) associated with the OpenGL rendering context by using the loop technique of 2., above**.**
6. Inside the loop:
   a. Create a window on a specific display by using the CreateWindow function. This function returns the window handle (HWND).
   b. Use GetDC to get a handle to the device context for the client area of a specific window, or for the entire screen (OR). Use the CreateDC function to create a device context (HDC) for the specified device.
   c. Use ChoosePixelFormat to match an appropriate pixel format supported by a device context to a given pixel format specification.
   d. Use SetPixelFormat to set the pixel format of the specified device context to the format specified.
   e. Use wglCreateContext to create a new OpenGL rendering context from device context (HDC).
   f. Use wglMakeCurrent to bind the GL context created in the above step as the current rendering context.
   g. Use clGetGLContextInfoKHR (See Section 9.7 of the *OpenCL Specification 1.1)* and CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR parameter to get the number of GL associated devices for CL context creation. If the number of devices is zero go to the next display in the loop. Otherwise, use clGetGLContextInfoKHR (See Section 9.7 of the *OpenCL Specification 1.1)* and the CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR parameter to get the device ID of the CL device associated with OpenGL context.
   h. Use clCreateContext (See Section 4.3 of the *OpenCL Specification 1.1)* to create the CL context (of type cl_context).

The following code demonstrates how to use WIN32 Windowing API in CL-GL interoperability on multi-GPU environment.

```
int xCoordinate = 0;
int yCoordinate = 0;

for (deviceNum = 0; EnumDisplayDevices(NULL,
                    deviceNum,
                    &dispDevice,
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
                    0); deviceNum++)
{

    if (dispDevice.StateFlags &
            DISPLAY_DEVICE_MIRRORING_DRIVER)
    {
        continue;
    }

    DEVMODE deviceMode;

    EnumDisplaySettings(dispDevice.DeviceName,
            ENUM_CURRENT_SETTINGS,
            &deviceMode);

    xCoordinate = deviceMode.dmPosition.x;
    yCoordinate = deviceMode.dmPosition.y;
    WNDCLASS windowclass;

    windowclass.style = CS_OWNDC;
    windowclass.lpfnWndProc = WndProc;
    windowclass.cbClsExtra = 0;
    windowclass.cbWndExtra = 0;
    windowclass.hInstance = NULL;
    windowclass.hIcon = LoadIcon(NULL, IDI_APPLICATION);
    windowclass.hCursor = LoadCursor(NULL, IDC_ARROW);
    windowclass.hbrBackground = (HBRUSH)GetStockObject(BLACK_BRUSH);
    windowclass.lpszMenuName = NULL;
    windowclass.lpszClassName = reinterpret_cast<LPCSTR>("SimpleGL");
               RegisterClass(&windowclass);
    gHwnd = CreateWindow(
    reinterpret_cast<LPCSTR>("SimpleGL"),
               reinterpret_cast<LPCSTR>(
               "OpenGL Texture Renderer"),
               WS_CAPTION | WS_POPUPWINDOW,
               xCoordinate,
               yCoordinate,
               screenWidth,
               screenHeight,
               NULL,
               NULL,
               windowclass.hInstance,
               NULL);
    hDC = GetDC(gHwnd);

    pfmt = ChoosePixelFormat(hDC, &pfd);

    ret = SetPixelFormat(hDC, pfmt, &pfd);

    hRC = wglCreateContext(hDC);

    ret = wglMakeCurrent(hDC, hRC);

    cl_context_properties properties[] =
    {
            CL_CONTEXT_PLATFORM,
            (cl_context_properties) platform,
            CL_GL_CONTEXT_KHR,
            (cl_context_properties) hRC,
            CL_WGL_HDC_KHR,
            (cl_context_properties) hDC,
            0
    };

    if (!clGetGLContextInfoKHR)
    {
        clGetGLContextInfoKHR = (clGetGLContextInfoKHR_fn)
        clGetExtensionFunctionAddress(
            "clGetGLContextInfoKHR");
    }
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
    size_t deviceSize = 0;
    status = clGetGLContextInfoKHR(properties,
                CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR,
                0,
                NULL,
                &deviceSize);

    if (deviceSize == 0)
    {
            // no interopable CL device found, cleanup
            wglMakeCurrent(NULL, NULL);
            wglDeleteContext(hRC);
            DeleteDC(hDC);
            hDC = NULL;
            hRC = NULL;
            DestroyWindow(gHwnd);
            // try the next display
            continue;
    }
    ShowWindow(gHwnd, SW_SHOW);
    //Found a winner
    break;
}

cl_context_properties properties[] =
{
    CL_CONTEXT_PLATFORM,
    (cl_context_properties) platform,
    CL_GL_CONTEXT_KHR,
    (cl_context_properties) hRC,
    CL_WGL_HDC_KHR,
    (cl_context_properties) hDC,
    0
};

status = clGetGLContextInfoKHR( properties,
            CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR,
            sizeof(cl_device_id),
            &interopDevice,
            NULL);

// Create OpenCL context from device's id
context = clCreateContext(properties,
            1,
            &interopDevice,
            0,
            0,
            &status);
```

E.1.3 Limitations

- It is recommended not to use GLUT in a multi-GPU environment.

E.2 Linux Operating System

E.2.1 Single GPU Environment

5.2.1.1 Creating CL Context from a GL Context

Using GLUT

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

1. Use `glutInit` to initialize the GLUT library and to negotiate a session with the windowing system. This function also processes the command-line options depending on the windowing system.
2. Use `glXGetCurrentContext` to get the current rendering context (GLXContext).
3. Use `glXGetCurrentDisplay` to get the display (Display *) that is associated with the current OpenGL rendering context of the calling thread.
4. Use `clGetGLContextInfoKHR` (see Section 9.7 of the *OpenCL Specification 1.1*) and the `CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR` parameter to get the device ID of the CL device associated with the OpenGL context.
5. Use `clCreateContext` (see Section 4.3 of the *OpenCL Specification 1.1*) to create the CL context (of type `cl_context`).

The following code snippet shows how to create an interoperability context using GLUT in Linux.

```
glutInit(&argc, argv);
glutInitDisplayMode(GLUT_RGBA | GLUT_DOUBLE);
glutInitWindowSize(WINDOW_WIDTH, WINDOW_HEIGHT);
glutCreateWindow("OpenCL SimpleGL");

gGLXContext glCtx = glXGetCurrentContext();

Cl_context_properties cpsGL[] =
{
    CL_CONTEXT_PLATFORM,

    (cl_context_properties)platform,

    CL_GLX_DISPLAY_KHR,

    (intptr_t) glXGetCurrentDisplay(),

    CL_GL_CONTEXT_KHR,

(   intptr_t) glCtx, 0};

status = clGetGLContextInfoKHR(cpsGL,
       CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR,
       sizeof(cl_device_id),
       &interopDevice,
       NULL);

// Create OpenCL context from device's id
context = clCreateContext(cpsGL,
           1,
           &interopDevice,
           0,
           0,
           &status);
```

Using X Window System

1. Use `XOpenDisplay` to open a connection to the server that controls a display.
2. Use `glXChooseFBConfig` to get a list of GLX frame buffer configurations that match the specified attributes.
3. Use glXChooseVisual to get a visual that matches specified attributes.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

4. Use `XCreateColormap` to create a color map of the specified visual type for the screen on which the specified window resides and returns the colormap ID associated with it. Note that the specified window is only used to determine the screen.
5. Use `XCreateWindow` to create an unmapped sub-window for a specified parent window, returns the window ID of the created window, and causes the X server to generate a CreateNotify event. The created window is placed on top in the stacking order with respect to siblings.
6. Use `XMapWindow` to map the window and all of its sub-windows that have had map requests. Mapping a window that has an unmapped ancestor does not display the window, but marks it as eligible for display when the ancestor becomes mapped. Such a window is called unviewable. When all its ancestors are mapped, the window becomes viewable and is visible on the screen if it is not obscured by another window.
7. Use `glXCreateContextAttribsARB` to initialize the context to the initial state defined by the OpenGL specification, and returns a handle to it. This handle can be used to render to any GLX surface.
8. Use `glXMakeCurrent` to make argrument3 (GLXContext) the current GLX rendering context of the calling thread, replacing the previously current context if there was one, and attaches argument3 (GLXcontext) to a GLX drawable, either a window or a GLX pixmap.
9. Use `clGetGLContextInfoKHR` to get the OpenCL-OpenGL interoperability device corresponding to the window created in step 5.
10. Use `clCreateContext` to create the context on the interoperable device obtained in step 9.

The following code snippet shows how to create a CL-GL interoperability context using the X Window system in Linux.

```
Display *displayName = XOpenDisplay(0);

int nelements;
GLXFBConfig *fbc = glXChooseFBConfig(displayName,
DefaultScreen(displayName), 0, &nelements);
      static int attributeList[] = { GLX_RGBA,
                                     GLX_DOUBLEBUFFER,
                                     GLX_RED_SIZE,
                                     1,
                                     GLX_GREEN_SIZE,
                                     1,
                                     GLX_BLUE_SIZE,
                                     1,
                                     None
                                    };
    XVisualInfo *vi = glXChooseVisual(displayName,
                          DefaultScreen(displayName),
                          attributeList);

    XSetWindowAttributes swa;
    swa.colormap = XCreateColormap(displayName,
                       RootWindow(displayName, vi->screen),
                       vi->visual,
                       AllocNone);
    swa.border_pixel = 0;
    swa.event_mask = StructureNotifyMask;
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
    Window win = XCreateWindow(displayName,
RootWindow(displayName, vi->screen),
                                10,
                                10,
                                WINDOW_WIDTH,
                                WINDOW_HEIGHT,
                                0,
                                vi->depth,
                                InputOutput,
                                vi->visual,
                                CWBorderPixel|CWColormap|CWEventMask,
                                &swa);

    XMapWindow (displayName, win);

    std::cout  << "glXCreateContextAttribsARB "
               << (void*) glXGetProcAddress((const
    GLubyte*)"glXCreateContextAttribsARB")
              << std::endl;

    GLXCREATECONTEXTATTRIBSARBPROC glXCreateContextAttribsARB =
    (GLXCREATECONTEXTATTRIBSARBPROC)
                            glXGetProcAddress((const
                            GLubyte*)"glXCreateContextAttribsARB");

    int attribs[] = {
                  GLX_CONTEXT_MAJOR_VERSION_ARB, 3,
                  GLX_CONTEXT_MINOR_VERSION_ARB, 0,
                  0
                  };

    GLXContext ctx = glXCreateContextAttribsARB(displayName,
                          *fbc,
                          0,
                          true,
                          attribs);
    glXMakeCurrent (displayName,

        win,

        ctx);

cl_context_properties cpsGL[] = {
CL_CONTEXT_PLATFORM,(cl_context_properties)platform,
CL_GLX_DISPLAY_KHR, (intptr_t) glXGetCurrentDisplay(),
CL_GL_CONTEXT_KHR, (intptr_t) gGlCtx, 0
                                 };
status = clGetGLContextInfoKHR( cpsGL,
                                     CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR,
                                     sizeof(cl_device_id),
                                     &interopDeviceId,
                                     NULL);

// Create OpenCL context from device's id
context = clCreateContext(cpsGL,
           1,
           &interopDeviceId,
           0,
           0,
           &status);
```

E.2.2 Multi-GPU Configuration

5.2.2.1 Creating CL Context from a GL Context

Using X Window System

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

1. Use `XOpenDisplay` to open a connection to the server that controls a display.
2. Use `ScreenCount` to get the number of available screens.
3. Use `XCloseDisplay` to close the connection to the X server for the display specified in the Display structure and destroy all windows, resource IDs (Window, Font, Pixmap, Colormap, Cursor, and GContext), or other resources that the client created on this display.
4. Use a FOR loop to enumerate the displays. To change the display, change the value of the environment variable DISPLAY.
5. Inside the loop:
   a. Use `putenv` to set the environment variable DISPLAY with respect to the display number.
   b. Use `OpenDisplay` to open a connection to the server that controls a display.
   c. Use `glXChooseFBConfig` to get a list of GLX frame buffer configurations that match the specified attributes.
   d. Use `glXChooseVisual` to get a visual that matches specified attributes.
   e. Use `XCreateColormap` to create a color map of the specified visual type for the screen on which the specified window resides and returns the colormap ID associated with it. Note that the specified window is only used to determine the screen.
   f. Use `XCreateWindow` to create an unmapped sub-window for a specified parent window, returns the window ID of the created window, and causes the X server to generate a CreateNotify event. The created window is placed on top in the stacking order with respect to siblings.
   g. Use `XMapWindow` to map the window and all of its sub-windows that have had map requests. Mapping a window that has an unmapped ancestor does not display the window but marks it as eligible for display when the ancestor becomes mapped. Such a window is called unviewable. When all its ancestors are mapped, the window becomes viewable and is visible on the screen, if it is not obscured by another window.
   h. Use `glXCreateContextAttribsARB` function to initialize the context to the initial state defined by the OpenGL specification and return a handle to it. This handle can be used to render to any GLX surface.
   i. Use `glXMakeCurrent` to make argrument3 (GLXContext) the current GLX rendering context of the calling thread, replacing the previously current context, if there was one, and to attach argument3 (GLXcontext) to a GLX drawable, either a window or a GLX pixmap.
   j. Use `clGetGLContextInfoKHR` to get the number of OpenCL-OpenGL interoperability devices corresponding to the window created in f, above.
   k. If the number of interoperable devices is zero, use `glXDestroyContext` to destroy the context created at step h, and go to step a; otherwise, exit from the loop (an OpenCL-OpenGL interoperable device has been found).

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

6. Use `clGetGLContextInfoKHR` to get the OpenCL-OpenGL interoperable device id.
7. Use `clCreateContext` to create the context on the interoperable device obtained in the previous step.

The following code segment shows how to create an OpenCL-OpenGL interoperability context on a system with multiple GPUs.

```
displayName = XOpenDisplay(NULL);
int screenNumber = ScreenCount(displayName);
XCloseDisplay(displayName);

for (int i = 0; i < screenNumber; i++)
{
    if (isDeviceIdEnabled())
    {
        if (i < deviceId)
        {
            continue;
        }
    }
    char disp[100];
    sprintf(disp, "DISPLAY=:0.%d", i);
    putenv(disp);
    displayName = XOpenDisplay(0);
    int nelements;
    GLXFBConfig *fbc = glXChooseFBConfig(displayName,
                           DefaultScreen(displayName),
                           0,
                           &nelements);
    static int attributeList[] = { GLX_RGBA,
                                   GLX_DOUBLEBUFFER,
                                   GLX_RED_SIZE,
                                   1,
                                   GLX_GREEN_SIZE,
                                   1,
                                   GLX_BLUE_SIZE,
                                   1,
                                   None
                                 };

    XVisualInfo *vi = glXChooseVisual(displayName,
                           DefaultScreen(displayName),
                           attributeList);
    XSetWindowAttributes swa;
    swa.colormap = XCreateColormap(displayName,
                       RootWindow(displayName, vi->screen),
                       vi->visual,
                       AllocNone);
    swa.border_pixel = 0;
    swa.event_mask = StructureNotifyMask;

    win = XCreateWindow(displayName,
                RootWindow(displayName, vi->screen),
                10,
                10,
                width,
                height,
                0,
                vi->depth,
                InputOutput,
                vi->visual,
                CWBorderPixel|CWColormap|CWEventMask,
                &swa);

    XMapWindow (displayName, win);
```


Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

```
    int attribs[] = {
                        GLX_CONTEXT_MAJOR_VERSION_ARB, 3,
                        GLX_CONTEXT_MINOR_VERSION_ARB, 0,
                        0
                    };

    GLXContext ctx = glXCreateContextAttribsARB(displayName,
                            *fbc,
                            0,
                            true,
                            attribs);
    glXMakeCurrent (displayName,
        win,
        ctx);

    gGlCtx = glXGetCurrentContext();
    properties cpsGL[] = {
    CL_CONTEXT_PLATFORM, (cl_context_properties)platform,
    CL_GLX_DISPLAY_KHR, (intptr_t) glXGetCurrentDisplay(),
    CL_GL_CONTEXT_KHR, (intptr_t) gGlCtx, 0
                          };

    size_t deviceSize = 0;
    status = clGetGLContextInfoKHR(cpsGL,
                        CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR,
                        0,
                        NULL,
                        &deviceSize);
    int numDevices = (deviceSize / sizeof(cl_device_id));


    if(numDevices == 0)
    {
        glXDestroyContext(glXGetCurrentDisplay(), gGlCtx);
        continue;
        }
    else
        {
            //Interoperable device found
            std::cout<<"Interoperable device found "<<std::endl;
            break;
         }
}

status = clGetGLContextInfoKHR( cpsGL,
                CL_CURRENT_DEVICE_FOR_GL_CONTEXT_KHR,
                sizeof(cl_device_id),
                &interopDeviceId,
                NULL);

// Create OpenCL context from device's id
context = clCreateContext(cpsGL,
                1,
                &interopDeviceId,
                0,
                0,
                &status);
```

E.3 Additional GL Formats Supported

The following is a list of GL formats beyond the minimum set listed in *The OpenCL Extension Specification, v 1.2* that AMD supports.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Table E.1 AMD-Supported GL Formats

GL internal format	CL images format
GL_ALPHA8	CL_A,CL_UNORM8
GL_R8,	CL_R, CL_UNORM_INT8
GL_R8UI	CL_R, CL_UNSIGNED_INT8
GL_R8I	CL_R, CL_SIGNED_INT8
GL_RG8	CL_RG, CL_UNORM_INT8
GL_RG8UI	CL_RG, CL_UNSIGNED_INT8
GL_RG8I	CL_RG, CL_SIGNED_INT8
GL_RGB8	CL_RGB, CL_UNORM_INT8
GL_RGB8UI	CL_RGB, CL_UNSIGNED_INT8
GL_RGB8I	CL_RGB, CL_SIGNED_INT8
GL_R16	CL_R, CL_UNORM_INT16
GL_R16UI	CL_R, CL_UNSIGNED_INT16
GL_R16I	CL_R, CL_SIGNED_INT16
GL_RG16	CL_RG, CL_UNORM_INT16
GL_RG16UI	CL_RG, CL_UNSIGNED_INT16
GL_RG16I	CL_RG, CL_SIGNED_INT16
GL_RGB16	CL_RGB, CL_UNORM_INT16
GL_RGB16UI	CL_RGB, CL_UNSIGNED_INT16
GL_RGB16I	CL_RGB, CL_SIGNED_INT16
GL_R32I	CL_R, CL_SIGNED_INT32
GL_R32UI	CL_R, CL_UNSIGNED_INT32
GL_R32F	CL_R, CL_FLOAT
GL_RG32I	CL_RG, CL_SIGNED_INT32
GL_RG32UI	CL_RG, CL_UNSIGNED_INT32
GL_RG32F	CL_RG, CL_FLOAT
GL_RGB32I	CL_RGB, CL_SIGNED_INT32
GL_RGB32UI	CL_RGB, CL_UNSIGNED_INT32
GL_RGB32F	CL_RGB, CL_FLOAT

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Appendix F
New and deprecated functions in OpenCL 2.0

F.1 New built-in functions in OpenCL 2.0

F.1.1 Work Item Functions

get_enqueued_local_size	local sizes in uniform part of NDRange
get_global_linear_id	unique 1D index for each work item in the NDRange
get_local_linear_id	unique 1D index for each work item in the work group

F.1.2 Integer functions

ctz	count trailing zero bits

F.1.3 Synchronization Functions

work_group_barrier	replaces barrier, adds scope

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

F.1.4 Address space quailfier functions

to_global	convert generic pointer to global pointer
to_local	convert genericpointer to local pointer
to_private	convert generic pointer to private pointer
get_fence	get fence appropriate to address space

F.1.5 Atomic functions

atomic_init	Initialize atomic value
atomic_work_item_fence	memory fence
atomic_store[_explicit]	atomic store
atomic_load[_explicit]	atomic load
atomic_exchange[_explicit]	atomic exchange
atomic_compare_exchange_strong[_explicit]	atomic compare and exchange (CAS)
atomic_compare_exchange_weak[_explicit]	atomic compare and exchange (CAS)
atomic_fetch_add[_explicit]	atomic fetch+add
atomic_fetch_sub[_explicit]	atomic fetch+sub
atomic_fetch_or[_explicit]	atomic fetch+or
atomic_fetcn_xor[_explicit]	atomic fetch+xor
atomic_fetch_and[_explicit]	atomic fetch+and
atomic_fetch_max_[explicit]	atomic fetch+max
atomic_fetch_min[_explicit]	atomic fetch+min
atomic_flag_test_and_set[_explicit]	atomic flag set
atomic_flag_clear[_explicit]	atomic flag clear

F.1.6 Image Read and Write Functions

read_imagef	Read from 2D depth [array] image
write_imagef	Write to 2D depth [array] image

F.1.7 Work group functions

work_group_all	Test all members of work group (and reduction)
work_group_any	Test any member of work group (or reduction)
work_group_broadcast	Brodcast value to every member of work group
work_group_reduce_add	Sum reduction across work group
work_group_reduce_max	Max reduction across work group
work_group_reduce_min	Min reduction across work group
work_group_scan_exclusive_add	Sum exclusive scan across work group
work_group_scan_exclusive_max	Max exclusive scan across work group

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

work_group_scan_exclusive_min	Min exclusive scan across work group
work_group_scan_inclusive_add	Sum inclusive scan across work group
work_group_scan_inclusive_max	Max inclusive scan across work group
work_group_scan_inclusive_min	Min inclusive scan across work group

F.1.8 Pipe functions

read_pipe	Read from pipe
write_pipe	Write to pipe
reserve_read_pipe	Reserve reads from pipe
reserve_write_pipe	Reserve writes to pipe
commit_read_pipe	Commit reserved pipe reads
commit_write_pipe	Commit reserved pipe writes
is_valid_reserve_id	Test reservation value
work_group_reserve_read_pipe	Work group read reservation
work_group_reserve_write_pipe	work group write reservation
work_group_commit_read_pipe	work group commit read reservation
work_group_commit_write_pipe	work group commit write reservation
get_pipe_num_packets	get current number of packets in pipe
get_pipe_max_packets	get capacity of pipe

F.1.9 Enqueueing Kernels

enqueue_kernel	Enqueue block as kernel
get_kernel_work_group_size	Query max work group size
get_kernel_preferred_work_group_size_multiple	Query preferred divisor of work group size
enqueue_marker	Enqueue a marker
retain_event	Increment refcount of event
release_event	Decrement refcount of event
create_user_event	Create user event
is_valid_event	Check if event is valid
set_user_event_status	Signal user event
capture_event_profiling_info	Schedule capture of profiling info
get_default_queue	Get default queue
ndrange_1D	Create 1D NDRange
ndrange_2D	Create 2D NDRange
ndrange_3D	Create 3D NDRange

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

F.1.10 Sub-groups

get_sub_group_size	Get size of current sub group
get_max_sub_group_size	Get size of largest sub group
get_num_sub_groups	Get number of sub groups in current work group
get_enqueued_num_sub_groups	Get number of sub groups in uniform work group
get_sub_group_id	Get id of current sub group
get_sub_group_local_id	Get ID of workitem in sub group
sub_group_barrier	Sub group barrier
sub_group_all	Test all members of work group (and reduction)
sub_group_any	Test any member of work group (or reduction)
sub_group_broadcast	Brodcast value to every member of work group
sub_group_reduce_add	Sum reduction across sub group
sub_group_reduce_max	Max reduction across sub group
sub_group_reduce_min	Min reduction across sub group
sub_group_scan_exclusive_add	Sum exclusive scan across sub group
sub_group_scan_exclusive_max	Max exclusive scan across sub group
sub_group_scan_exclusive_min	Min exclusive scan across sub group
sub_group_scan_inclusive_add	Sum inclusive scan across sub group
sub_group_scan_inclusive_max	Max inclusive scan across sub group
sub_group_scan_inclusive_min	Min inclusive scan across sub group
sub_group_reserve_read_pipe	sub group read reservation
sub_group_reserve_write_pipe	sub group write reservation
sub_group_commit_read_pipe	sub group commit read reservation
sub_group_commit_write_pipe	sub group commit write reservation
get_kernel_sub_group_count_for_ndrange	number of sub groups in uniform part of NDRange
get_kernel_max_sub_group_size_for_ndrange	maximum sub group size for a block

F.2 Deprecated built-ins

barrier

mem_fence

read_mem_fence

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

write_mem_fence
atomic_add
atomic_sub
atomic_xchg
atomic_inc
atomic_dec
atomic_cmpxchg
atomic_min
atomic_max
atomic_and
atomic_or
atomic_xor

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

F.3 New runtime APIs in OpenCL 2.0

F.3.1 New Types

cl_device_svm_capabilities	Returned by clGetDeviceInfo(...CL_DEVICE_SVM_CAPABILITIES...)
cl_queue_properties	See clCreateCommandQueueWithProperties
cl_svm_mem_flags	See clSVMAlloc
cl_pipe_properties	See clCreatePipe
cl_pipe_info	See clGetPipeInfo
cl_sampler_properties	See clCreateSamplerWithProperties
cl_kernel_exec_info	See clSetKernelExecInfo
cl_image_desc	A field name changed from buffer to mem_object
cl_kernel_sub_group_info	See clGetKernelSubGroupInfoKHR

F.3.2 New Macros

CL_INVALID_PIPE_SIZE
CL_INVALID_DEVICE_QUEUE
CL_VERSION_2_0
CL_DEVICE_QUEUE_ON_HOST_PROPERTIES
CL_DEVICE_MAX_READ_WRITE_IMAGE_ARGS
CL_DEVICE_MAX_GLOBAL_VARIABLE_SIZE
CL_DEVICE_QUEUE_ON_DEVICE_PROPERTIES
CL_DEVICE_QUEUE_ON_DEVICE_PREFERRED_SIZE
CL_DEVICE_QUEUE_ON_DEVICE_MAX_SIZE
CL_DEVICE_MAX_ON_DEVICE_QUEUES
CL_DEVICE_MAX_ON_DEVICE_EVENTS
CL_DEVICE_SVM_CAPABILITIES
CL_DEVICE_GLOBAL_VARIABLE_PREFERRED_TOTAL_SIZE
CL_DEVICE_MAX_PIPE_ARGS
CL_DEVICE_PIPE_MAX_ACTIVE_RESERVATIONS
CL_DEVICE_PIPE_MAX_PACKET_SIZE
CL_DEVICE_PREFERRED_PLATFORM_ATOMIC_ALIGNMENT
CL_DEVICE_PREFERRED_GLOBAL_ATOMIC_ALIGNMENT
CL_DEVICE_PREFERRED_LOCAL_ATOMIC_ALIGNMENT
CL_QUEUE_ON_DEVICE
CL_QUEUE_ON_DEVICE_DEFAULT
CL_DEVICE_SVM_COARSE_GRAIN_BUFFER
CL_DEVICE_SVM_FINE_GRAIN_BUFFER
CL_DEVICE_SVM_FINE_GRAIN_SYSTEM

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

CL_DEVICE_SVM_ATOMICS
CL_QUEUE_SIZE
CL_MEM_SVM_FINE_GRAIN_BUFFER
CL_MEM_SVM_ATOMICS
CL_sRGB
CL_sRGBx
CL_sRGBA
CL_sBGRA
CL_ABGR
CL_MEM_OBJECT_PIPE
CL_MEM_USES_SVM_POINTER
CL_PIPE_PACKET_SIZE
CL_PIPE_MAX_PACKETS
CL_SAMPLER_MIP_FILTER_MODE
CL_SAMPLER_LOD_MIN
CL_SAMPLER_LOD_MAX
CL_PROGRAM_BUILD_GLOBAL_VARIABLE_TOTAL_SIZE
CL_KERNEL_ARG_TYPE_PIPE
CL_KERNEL_EXEC_INFO_SVM_PTRS
CL_KERNEL_EXEC_INFO_SVM_FINE_GRAIN_SYSTEM
CL_COMMAND_SVM_FREE
CL_COMMAND_SVM_MEMCPY
CL_COMMAND_SVM_MEMFILL
CL_COMMAND_SVM_MAP
CL_COMMAND_SVM_UNMAP
CL_PROFILING_COMMAND_COMPLETE

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

F.3.3 New API calls

clCreateCommandQueueWithProperties	See section 5.1
clCreatePipe	See section 5.4.1
clGetPipeInfo	See section 5.4.2
clSVMAlloc	See section 5.6.1
clSVMFree	See section 5.6.1
clEnqueueSVMFree	See section 5.6.1
clEnqueueSVMMemcpy	See section 5.6.1
clEnqueueSVMMemFill	See section 5.6.1
clEnqueueSVMMap	See section 5.6.1
clEnqueueSVMUnmap	See section 5.6.1
clCreateSamplerWithProperties	See section 5.7.1
clSetKernelArgSVMPointer	See section 5.9.2
clSetKernelExecInfo	See section 5.9.2
clGetKernelSubGroupInfoKHR	See section 9.17.2.1

F.4 Deprecated runtimes

clCreateCommandQueue
clCreateSampler
clEnqueueTask

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Index

Symbols

_cdecl calling convention
  Linux . . . 7
_global atomics . . . 19
_local atomics . . . 19
_stdcall calling convention
  Windows . . . 7
.amdil
  generating . . . 3
.comment
  BIF binary . . . 1
  storing OpenCL and driver versions that created the binary . . . 1
.llvmir
  generating . . . 3
  storing OpenCL immediate representation (LLVM IR) . . . 1
.rodata
  storing OpenCL runtime control data . . . 1
.shstrtab
  forming an ELF . . . 1
.source
  storing OpenCL source program . . . 1
.strtab
  forming an ELF . . . 1
.symtab
  forming an ELF . . . 1
.text
  generating . . . 3
  storing the executable . . . 1

Numerics

1D address . . . 10
2D
  address . . . 10
2D addresses
  reading and writing . . . 10

A

access
  memory . . . 5, 9
accumulation operations
  NDRange . . . 1
address
  1D . . . 10
  2D . . . 10
  normalized . . . 10
  un-normalized . . . 10
allocating
  images
    OpenCL . . . 4
  memory
    selecting a device . . . 4
  memory buffer
    OpenCL program model . . . 4
ALUs
  arrangement of . . . 8
AMD Accelerated Parallel Processing
  implementation of OpenCL . . . 1
  open platform strategy . . . 1
  programming model . . . 2
  relationship of components . . . 1
  software stack . . . 1
    Installable Client Driver (ICD) . . . 1
AMD APP KernelAnalyzer . . . 1
AMD Core Math Library (ACML) . . . 1
AMD GPU
  number of compute units . . . 8
AMD Radeon HD 68XX . . . 15
AMD Radeon HD 69XX . . . 15
AMD Radeon HD 75XX . . . 15
AMD Radeon HD 77XX . . . 15
AMD Radeon HD 78XX . . . 15
AMD Radeon HD 79XX series . . . 15
AMD Radeon HD 7XXX . . . 2
AMD Radeon R9 290X . . . 4, 8
AMD supplemental compiler . . . 6
  -g option . . . 6
AMD supplemental compiler option
  -f[n-]bin-source . . . 6
  -f[no-]bin-amdil . . . 6
  -f[no-]bin-exe . . . 6
  -f[no-]bin-llvmir . . . 6
AMD vendor-specific extensions . . . 4
amd_bitalign

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

built-in function 8
amd_bytealign
built-in function 8
amd_lerp
built-in function 8
AMD_OCL_BUILD_OPTIONS
environment variables 5
AMD_OCL_BUILD_OPTIONS_APPEND
environment variable 5
amd_pack
built-in function 7
amd_sad
buillt-in function 8
amd_sad4
built-in function 8
amd_sadhi
built-in function 9
amd_unpack0
built-in function 7
amd_unpack1
built-in function 7
amd_unpack2
built-in function 7
amd_unpack3
built-in function 7
API
C++ 2
naming extension functions 1
OpenCL 1
platform
querying 1
processing calls 9
API commands
three categories 6
application code
developing Visual Studio 4
application kernels
device-specific binaries 2
arrangement of ALUs 8
atomics
_global 19
_local 19

B

barrier
command-queue 4
barriers
execution order 4
work-group 4
work-items
encountering 4
BIF
.comment
storing OpenCL and driver versions that created the binary 1
.llvmir
storing immediate representation (LLVM IR) 1
.source
storing OpenCL source program 1
binary
.comment section 1
bitness 3
changing default behavior 1
ELF special sections 1
options to control what is contained in the binary 3
overview 1
binary
application kernels 2
controlling
BIF options 3
CPU 1
generating
in OpenCL 1
LLVM AS 1
GPU 1
Binary Image Format (BIF)
See BIF
bitness
BIF 3
branch
granularity
work-item 4
instructions 4
branching
flow control 4
breakpoint
CL kernel function 3
host code 3
no breakpoint is set 2
setting 2, 4
sample GDB debugging session 3
setting a 2
buffer
command queue 9
global 9
source or destination for instruction . . . 10
storing writes to random memory locations 9
relationship
sample code 4
build log
printing out 16
built-in function

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

amd_bitalign 8
amd_bytealign 8
amd_lerp 8
amd_pack 7
amd_sad 8
amd_sad4 8
amd_sadhi 9
amd_unpack0 7
amd_unpack1 7
amd_unpack2 7
amd_unpack3 7
built-in functions
for OpenCL language
cl_amd_media_ops 7, 9
OpenCL C programs
cl_amd_printf 12
variadic arguments 12
writing output to the stdout stream 12
burst write 10

C

C front-end
compiler 2
C kernels
predefined macros 13
C program sample
OpenCL 10
C programming
OpenCL 3
C++ API 2
C++ bindings
OpenCL programming 14
C++ extension
unsupported features 2
C++ files
compiling 3
C++ kermel language . . . . . iii, 1
C++ kernels
building 3
C++ templates 5
cache
L1 10
L2 10
texture system 10
call error
OpenCL 16
calling convention
Linux 7
Windows 7
character extensions 1
searching for substrings 2
character sequence
format string 12
CL context
associate with GL context 1
CL kernel function
breakpoint 3
CL options
change during runtime 5
cl_amd_device_attribute_query extension
querying AMD-specific device attributes . . 5
cl_amd_event_callback extension
registering event callbacks for states 6
cl_amd_fp64 extension 4
cl_amd_media_ops extension
adding built-in functions to OpenCL language
7, 9
cl_amd_printf extension 12
cl_ext extensions 4
cl_khr_fp64
supported function 15
classes
passing between host and device 3
clBuildProgram
debugging OpenCL program 2
clCreateKernel
C++extension 2
clEnqueue commands 5
clEnqueueNDRangeKernel
setting breakpoint in the host code 3
clGetPlatformIDs() function
available OpenCL implementations 1
clGetPlatformInfo() function
available OpenCL implementations 1
querying supported extensions for OpenCL
platform 1
C-like language
OpenCL 3
code
basic programming steps 10
ICD-compliant version 1, 2
parallel min() function 19
pre-ICD snippet 1, 2
running
on Linux 7
on Windows 6
runtime steps 19
code requirements
Installable Client Driver (ICD) 1
CodeXL GPU Debugger 1
command processor
transfer from system to GPU 9
command processors
concurrent processing of command queues 7
command queue 4
associated with single device 4

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

barrier
enforce ordering within a single queue . 4
creating device-specific. 4
elements
constants. 9
kernel execution calls 9
kernels 9
transfers between device and host 9
executing kernels 4
execution. 9
moving data 4
no limit of the number pointing to the same device 4
OpenCL. 4
command queues 7
multiple 6
command-queue barrier 4
commands
API
three categories 6
buffer. 9
clEnqueue 5
driver layer issuing 9
driver layer translating 9
event 6
GDB 3
kernel 6
memory 6
OpenCL API functions 5
queue 9
communication and data transfers between system and GPU
PCIe 8
communication between the host (CPU) and the GPU 8
compilation
error
kernel 16
compile time
resolving format string 12
compiler
LLVM framework. 2
set to ignore all extensions 2
toolchain 1
back-end 1
OpenCL. 1
sharing front-end 1
sharing high-level transformations 1
transformations 1
using standard C front-end. 2
compiler option
-f[no-]bin-amdil 4
-f[no-]bin-exe 4
-f[no-]bin-llvmir 4
-f[no-]bin-source 4
-g 4
-O0 4
-save-temps 4
compiling
an OpenCL application 2
C++ files 3
kernels. 1
on Linux
building 32-bit object files on a 64-bit system 3
linking to a 32-bit library. 3
linking to a 64-bit library. 3
OpenCL on Linux 3
OpenCL on Windows 2
Intel C (C++) compiler 2
setting project properties 2
Visual Studio 2008 Professional Edition 2
the host program 2
computation
data-parallel model 2
compute device structure
GPU 6, 2
compute kernel
data-parallel granularity 2
definition 1
strengths
computationally intensive applications . . 1
wavefronts. 2
workgroups 2
compute unit
mapping 2
stream cores. 3
compute unites
number in AMD GPU 8
compute units
290X devices 6
independent operation 4
number in AMD GPUs 8
structured in AMD GPUs 8
constants
caching 10
command queue elements 9
constraints
of the current LDS model. 2
context
relationship
sample code 4
contexts
associating CL and GL. 1
copying data
implicit and explicit 6

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

copying processes 7
CPU
  binaries 1
  code
    parallel min() function 19
  communication between host and GPU . . 8
  predefined macros 13
  processing 2
    LLVM AS 2
    OpenCL runtime 1
  skip copying between host memory and PCIe memory 7
Creating CL context
  from a GL Context 10
cygwin
  GDB running 4
cygwin/minGW 1

D

-D name
  OpenCL supported options 4
data
  computations
    select a device 4
  fetch units 9
  moving using corresponding command queue 4
  parallelism
    grouping 2
  transfers
    select a device 4
    to the optimizer 2
dataflow between host (CPU) and GPU 6
data-parallel granularity
  compute kernels 2
data-parallel programming model
  executing non-graphic functions 2
debug
  information
    creating metadata structures 2
debugger
  CodeXL GPU 1
  GNU project 4, 1
  kernel symbols not visible 2
debugging
  kernels 2
  OpenCL 1
    kernels in Windows 4
  OpenCL programs 1
  session
    GDB sample 3
    hello world kernel sample 2
  setting the environment 2
definition
  kernel 1
  NDRange 1
  wavefront 4
derived classes 3
device
  extension support listing 15
  fission extension support in OpenCL 4
  list
    function call query 2
  no limit of number of command queues . . 4
  relationship
    sample code 4
device-optimal access pattern
  threading 19
devices
  R9 290X series 8
device-specific
  operations
    kernel execution 4
    program compilation 4
Direct Memory Access (DMA)
  engine 9
    signaling transfer is completed 9
  transfers 9
    parallelization 9
directives
  extension name overrides 2
  order 2
domains
  of synchronization 4
    command-queue 4
    work-items 4
double copying
  memory bandwidth 7
double-precision floating-point
  performing operations 3, 4
driver layer
  issuing commands 9
  translating commands 9

E

element
  work-item 3
ELF
  .rodata
    storing OpenCL runtime control data . . 1
  .shstrtab
    forming an ELF 1
  .strtab
    forming an ELF 1
  .symtab
    forming an ELF 1

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

.text
storing the executable 1
format 1
forming 1
header fields 2
special sections
BIF 1
enforce ordering
between or within queues
events 5
synchronizing a given event 3
within a single queue
command-queue barrier 4
engine
DMA 9
enqueuing
commands in OpenCL 5
multiple tasks
parallelism 2
native kernels
parallelism 2
environment variable
AMD_OCL_BUILD_OPTIONS 5
AMD_OCL_BUILD_OPTIONS_APPEND . . 5
setting to avoid source changes 2
event
commands 6
enforces ordering
between queues 5
within queues 5
synchronizing 3
event commands 6
events
forced ordering between 5
exceptions
C++ 6
executing
branch 4
kernels 2, 3
using corresponding command queue . . 4
kernels for specific devices
OpenCL programming model 3
loop 5
non-graphic function
data-parallel programming model 2
execution
command queue 9
of a single instruction over all work-items . 2
OpenCL applications 6
order
barriers 4
single stream core 10
explicit copying of data 6
extension
cl_amd_popcnt 7
clCreateKernel 2
extension function pointers 3
extension functions
NULL and non-Null return values 3
extension support by device
for devices 1 15
for devices 2 and CPUs 16
extensions
all 2
AMD vendor-specific 4
approved by Khronos Group 1
approved list from Khronos Group 3
character strings 1
cl_amd_device_attribute_query 5
cl_amd_event_callback
registering event callbacks for states . . . 6
cl_amd_fp64 4
cl_amd_media_ops 7, 9
cl_amd_printf 12
cl_ext 4
compiler set to ignore 2
device fission 4
disabling 2
enabling 2, 3
FunctionName string 3
kernel code compilation
adding defined macro 3
naming conventions 1
optional 1
provided by a specific vendor 1
provided collectively by multiple vendors . . 1
querying for a platform 1
querying in OpenCL 1
same name overrides 2
use in kernel programs 2

F

-f[n-]bin-source
AMD supplemental compiler option 6
-f[no-]bin-amdil
AMD supplemental compiler option 6
compiler option 4
-f[no-]bin-exe
AMD supplemental compiler option 6
compiler option 4
-f[no-]bin-llvmir
AMD supplemental compiler option 6
compiler option 4
-f[no-]bin-source
compiler option 4
fetch unit

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

loads . . . 9
processing . . . 10
stores . . . 9
streaming stores . . . 9
transferring the work-item . . . 9
fetches
memory
stalls . . . 11
floating point operations
double-precision . . . 4
single-precision . . . 4
flow control . . . 4
branching . . . 4
execution of a single instruction over all work-items . . . 2
forced ordering of events . . . 5
format string . . . 12
conversion guidelines . . . 12
resolving compile time . . . 12
front-end
ensuring the kernels meet OpenCL specification . . . 2
performs
semantic checks . . . 2
syntactic checks . . . 2
standard C . . . 2
translating . . . 2
front-end supports
additional data-types
float8 . . . 2
int4 . . . 2
additional keywords
global . . . 2
kernel . . . 2
built-in functions
barrier() . . . 2
get_global_id() . . . 2
function call
querying . . . 2
function names
undecorated in Windows . . . 7
FunctionName string
address of extension . . . 3

G

-g
compiler option . . . 4
experimental feature . . . 6
-g option
passing to the compiler . . . 2
gather/scatter model
Local Data Store (LDS) . . . 2
gcc
not supported . . . 4
GCC option syntax . . . 3
GDB
GNU project debugger . . . 4, 1
sample session . . . 3
GDB (GNU project debugger) . . . 6
command . . . 3
documentation . . . 4
running cygwin . . . 4
running minGW . . . 4
generate binary images offline
cl_amd_offline_devices . . . 6
GL context
associate with CL context . . . 1
global buffer . . . 9, 10
source or destination for instruction . . . 10
storing writes to random memory locations 9
global memory (VRAM) . . . 9
global scope . . . 4
global work-size
defined . . . 3
dividing into work-groups . . . 3
GLUT windowing system . . . 2
GNU project debugger
GDB, description . . . 2
GNU project debugger (GDB)
See GDB
GPU . . . 1
binaries . . . 1
communication between host and GPU . . . 8
copying data from host to GPU . . . 6
dataflow between host and GPU . . . 6
parallel min() function code . . . 19
parallelizing large numbers of work-items 10
predefined macros . . . 13
processing . . . 2
LLVM IR-to_CAL IL module . . . 2
OpenCL runtime . . . 1
programming . . . 2
scheduling . . . 10
specific macros . . . 1
storing writes to random memory locations 9
structure . . . 6, 2
GPUs
dedicating for specific operations . . . 10
masking . . . 10
granularity
branch . . . 4
data-parallel . . . 2

H

hardware
overview . . . 1

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Hawaii
see R9 290X series devices or AMD Radeon R9 290X . . . 8
header fields
in ELF . . . 2
hello world sample kernel . . . 2
hierarchical subdivision
OpenCL data-parallel programming model 2
host
communication between host and GPU . . . 8
copying data from host to GPU . . . 6
dataflow between host and GPU . . . 6
program
OpenCL . . . 2
program compiling . . . 2
host code
breakpoint . . . 3
platform vendor string . . . 3
setting breakpoint . . . 3
clEnqueueNDRangeKernel . . . 3
host/device architecture single platform
consisting of a GPU and CPU . . . 3

I

-I dir
OpenCL supported options . . . 4
idle stream cores . . . 4
IL
complete . . . 1
incomplete . . . 1
image
reads . . . 10
implicit copying of data . . . 6
index space
n-dimensional . . . 2
inheritance
strict and multiple . . . 1
input stream
NDRange . . . 1
Installable Client Driver (ICD) . . . 1
AMD Accelerated Parallel Processing
software stack . . . 1
compliant version of code . . . 1, 2
overview . . . 1
pre-ICD code snippet . . . 1, 2
using . . . 1
instruction
branch . . . 4
global buffer . . . 10
kernel . . . 9
sequence
stream cores . . . 3
instructions
scalar and vector . . . 7
integer
performing operations . . . 3, 4
Intel C (C++) compiler
compiling OpenCL on Windows . . . 2
interoperability context
code for creating . . . 3
interrelationship of memory domains . . . 6

K

kernel
code
parallel min() function . . . 20
code compilation
adding a defined macro with the name of the extension . . . 3
code example . . . 13
command queue elements . . . 9
commands . . . 6
compilation error . . . 16
compiling . . . 1
compute
definition . . . 1
strengths . . . 1
creating within programs . . . 4
debugging . . . 2
definition of . . . 1
device-specific binaries . . . 2
distributing in OpenCL . . . 1
executed as a function of multi-dimensional
domains of indices . . . 3
executing . . . 2
using corresponding command queue . . . 4
execution
device-specific operations . . . 4
execution calls
command queue elements . . . 9
hello world sample . . . 2
instructions over PCIe bus . . . 9
keyword . . . 1
loading . . . 2
no breakpoint set . . . 2
overloading . . . 2
program
OpenCL . . . 2
programming
enabling extensions . . . 3
programs using extensions . . . 2
running on compute unit . . . 2
stream . . . 10
submitting for execution . . . 4
synchronization points . . . 2
work-item . . . 2

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

kernel and function overloading 1
kernel commands 6
kernel_name
construction 2
kernels
debugging 2
kernels and shaders 2
Khronos
approved list of extensions 3
website 14

L

L1 cache 10
L2 cache 10, 7
latency
hiding in memory 10
latency hiding 4, 11
launching
threads 19
LDS
description 2
gather/scatter model 2
size allocated to work-group 2
using local memory 9
LDS model constraints 2
data sharing 2
memory accesses outside the work-group 2
size is allocated per work-group 2
library
SDKUtil 3
linking
creating an executable 3
in the built-in OpenCL functions 2
object files 3
OpenCL on Linux 3
options
SDKUtil library 3
to a 32-bit library compiling on Linux 3
to a 64-bit library compiling on Linux 3
Linux
building 32-bit object files on a 64-bit system 3
calling convention
_cdecl 7
compiling OpenCL 3
linking 3
linking
to a 32-bit library 3
to a 64-bit library 3
linking options
SDKUtil library 3
running code 7
SDKUtil library 3
list of supported extensions
approved by the Khronos Group 3
LLVM
compiler 2
framework
compiler 2
linker 2
LLVM AS
CPU processing 2
generating binaries 1
LLVM IR
BIF 1
compatibility 1
enabling recompilation to the target 1
front-end translation 2
generating a new code 1
LLVM IR-to-CAL IL module
GPU processing 2
Local Data Store (LDS)
See LDS
loop
execute 5
Low-Level Virtual Machine (LLVM)
See LLVM

M

macros
GPU-specific 1
predefined
CPU 13
GPU 13
OpenCL C kernels 13
mapping
executions onto compute units 2
OpenCL 3
work-items onto n-dimensional grid (ND-Range) 3
work-items to stream cores 2
masking GPUs 10
mem_fence operation 4
memories
interrelationship of 5
memory
access 5, 9
allocation
select a device 4
architecture 5
bandwidth 7
double copying 7
commands 6
domains
interrelationship 6

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

fence
barriers . . . 3
operations . . . 3
global (VRAM) . . . 9
hiding latency . . . 10, 11
loads . . . 9
object allocation
OpenCL context . . . 3
OpenCL domains . . . 5
read operations . . . 10
request . . . 9
stores . . . 9
streaming . . . 9
system pinned . . . 7
transfer management . . . 7
write operations . . . 10
memory access
stream cores . . . 9
memory commands . . . 6
metadata structures
holding debug information . . . 2
OpenCL-specific information . . . 2
minGW
GDB running . . . 4
minGW/cygwin . . . 1
multi-GPU environment
use of GLUT . . . 7

N

namespaces
C++ support for . . . 4
supported feature in C++ . . . 1
naming conventions
API extension functions . . . 1
elements contained in extensions . . . 1
enumerated values . . . 1
extensions . . . 1
Khronos Group approved . . . 1
provided by a specific vendor . . . 1
provided by multiple vendors . . . 1
n-dimensional grid (ND-Range) . . . 3
n-dimensional index space
NDRange . . . 2
NDRange
accumulation operations . . . 1
defnition . . . 1
input streams . . . 1
n-dimensional index space . . . 2
random-access functionality . . . 1
reduction operations . . . 1
variable output counts . . . 1
non-active work-items . . . 4
normalized addresses . . . 10
NULL and non-Null return values
extension functions . . . 3

O

-O0
compiler option . . . 4
object files
linking . . . 3
open platform strategy
AMD Accelerated Parallel Processing . . . 1
OpenCL
Accelerated Parallel Processing
implementation . . . 1
adding built-in functions to the language
cl_amd_media_ops extension . . . 7
allocating images . . . 4
API . . . 1
applications
execution . . . 6
Binary Image Format (BIF)
overview . . . 1
building
create a context . . . 3
programs . . . 1
querying the runtime . . . 3
the application . . . 3
C printf . . . 12
C programming . . . 3
call error . . . 16
checking for known symbols . . . 2
C-like language with extensions
for parallel programming . . . 3
coding . . . 1
compiler
toolchain . . . 1
compiler and runtime components . . . 1
compiler options
-D name . . . 4
-I dir . . . 4
compiling
on Linux . . . 3
linking . . . 3
on Windows . . . 2
the program . . . 2
context
memory object allocation . . . 3
conversion guidelines
format string . . . 12
CPU processing . . . 1
create kernels within programs . . . 4
create one or more command queues . . . 4
create programs to run on one or more
devices . . . 4

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

creating a context
selecting a device 4
data-parallel model
hierarchical subdivision 2
debugging 1
clBuildProgram 2
desired platform 3
selection 3
directives to enable or disable extensions. 2
distributing the kernel 1
enqueued commands 5
ensuring the kernels meet specification . . 2
extensions
enabling or disabling 2
following same pattern 4
functions 2
general compilation path of applications . . 1
generating
.amdil 3
.llvmir 3
.text 3
a binary 1
GPU processing 1
host program 2
ICD code requirements 1
implementations
use clGetPlatformIds() function 1
use clGetPlatformInfo() function 1
Installable Client Driver (ICD) 1
introductory sample
C++ bindings 14
kernel compiling 1
kernel symbols
not visible in debugger. 2
list
of available implementations 1
of commands 5
mapping 3
memory domains 5
metadata structures 2
minimalist C program sample 10
optional
extensions 1
kernel program 2
performance
libraries components 1
profiling components 1
printf capabilities. 12
programmers introductory sample 14
programming model 3
allocating memory buffers 4
executing kernels for specific devices. . 3
queues of commands 3
reading/writing data 3
providing an event 3
querying
extensions 1
supported extensions using clGetPlat-
formInfo() 1
read data back to the host from device . . 4
recompiling LLVM IR to generate a new code
1
re-written applications become source . . . 1
running
data-parallel work. 3
programs 1
task-parallel work 3
runtime 1
changing options 6
interface 2
post-processes the incomplete AMD IL
from OpenCL compiler 1
using LLVM AS 1
setting breakpoint 2
settings for compiling on Windows 2
storing immediate representation (LLVM IR)
.llvmir 1
storing OpenCL and driver versions
.comment 1
storing source program
.source 1
submit the kernel for execution 4
supported standard
compiler options 4
synchronizing a given event 3
write data to device 4
OpenCL device
general overview 6
OpenCL programs
debugging 1
operation
mem_fence 4
operations
device-specific
kernel execution 4
program compilation. 4
double-precision floating point 4
integer. 4
memory-read 10
memory-write 10
single-precision floating point. 4
optimizer

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

transfer of data 2
optional extensions
for OpenCL 1
overloading
in C++ language 4
kernel 2
kernel and function 1
overview
software and hardware 1

P

parallel min() function
code sample 21
example programs 19
kernel code 20
programming techniques 19
runtime code 19
steps 19
parallel programming
memory fence
barriers 3
operations 3
parallelism
enqueuing
multiple tasks 2
native kernels 2
using vector data types 2
parallelization
DMA transfers 9
GPU 10
passing a class between host to the device . 6
path to the kernel file
relative to executable 6, 7
PCIe
communication between system and GPU 8
data transfers between system and GPU . 8
kernel instructions 9
overview 8
skip copying between host memory and PCIe memory 7
throughput 8
performance
work-groups 2
pipelining 7
platform vendor string
remains constant for a particular vendor's implementation 3
searching for desired OpenCL platform . . . 3
vs platform name string 3
point (barrier) in the code 2
population count
extension 7
post-processing incomplete AMD IL 1
pre-ICD code snippet 1, 2
processing by command processors 7
processing elements
SIMD arrays 8
program
examples 10
simple buffer write 10
programming
basic steps with minimum code 10
GPU 2
techniques
simple tests
parallel min() function 19
programming model
AMD Accelerated Parallel Processing 2
OpenCL 3
executing kernels for specific devices . . 3
queues of commands 3
reading/writing data 3
project property settings
compiling on Windows 2

Q

querying
AMD-specific device attributes 5
extensions
for a list of devices 2
for a platform 1
OpenCL 1
for a specific device 2
for available platforms 2
the platform API 1
the runtime
OpenCL building 3
queue
command 4, 9

R

R9 290X series devices 8
random memory location
GPU storage of writes 9
random-access functionality
NDRange 1
read imaging 10
read-only buffers
constants 10
running code
on Linux 7
on Windows 6
runtime
change CL options 5
code
parallel min() function 19

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

interface
OpenCL . . . 2
OpenCL . . . 1
changing options . . . 6
system functions . . . 5

S

sample code
relationship between
buffer(s) . . . 4
command queue(s) . . . 4
context(s) . . . 4
device(s) . . . 4
kernel(s) . . . 4
relationship between context(s) . . . 4
-save-temps
compiler option . . . 4
SAXPY function
code sample . . . 16
SC cache . . . 7
scalar instructions . . . 7
scalar unit data cache
SC cache . . . 7
scalra instructions . . . 7
scheduling
GPU . . . 10
work-items
for execution . . . 3
range . . . 2
scope
global . . . 4
SDKUtil library . . . 3
linking options . . . 3
Linux . . . 3
set
a breakpoint . . . 2
shader architecture
unified . . . 2
shaders and kernels . . . 2
SIMD arrays
processing elements . . . 8
simple buffer write
code sample . . . 12
example programs . . . 10
simple testing
programming techniques
parallel min function . . . 19
single device associated with command queue 4
single stream core execution . . . 10
single-precision floating-point
performing operations . . . 3, 4
software
overview . . . 1
spawn order
of work-item . . . 1
sequential . . . 1
stalls
memory fetch request . . . 11
static C++ kernel language . . . iii, 1
stdout stream
writing output associated with the host application . . . 12
stream core
compute units . . . 3
executing kernels . . . 3
idle . . . 4
instruction sequence . . . 3
processing elements . . . 3
stall . . . 11
due to data dependency . . . 12
stream kernel . . . 10
supplemental compiler options . . . 6
synchronization
command-queue barrier . . . 4
domains . . . 4
command-queue . . . 4
work-items . . . 4
events . . . 4
points
in a kernel . . . 2
synchronizing
a given event . . . 3
event
enforce the correct order of execution . 3
through barrier operations
work-items . . . 3
through fence operations
work-items . . . 3
syntax
GCC option . . . 3
system
pinned memory . . . 7

T

templates
C++ . . . 5
kernel, member, default argument, limited class, partial . . . 1
terminology . . . 1
texture system
caching . . . 10
thread
launching . . . 19
threading
device-optimal access pattern . . . 19

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

throughput
PCIe 8
timing
of simplified execution of work-items
single stream core 10
toolchain
compiler 1
transcendental
core 4
performing operations 3
transfer
between device and host
command queue elements 9
data
select a device 4
to the optimizer 2
DMA 9
from system to GPU
command processor 9
DMA engine 9
management
memory 7
work-item to fetch unit 9

U

unified shader architecture 2
un-normalized addresses 10

V

variable output counts
NDRange 1
variadic arguments
use of in the built-in printf 12
vector data types
parallelism 2
vector instructions 7
vendor
platform vendor string 3
vendor name
matching platform vendor string 3
vendor-specific extensions
AMD 4
Very Long Instruction Word (VLIW)
instruction 4
work-item 4
Visual Studio 2008 Professional Edition 4
compiling OpenCL on Windows 2
developing application code 4
VRAM
global memory 9

W

wavefront
block of work-items 4
combining paths 4
concept relating to compute kernels 2
definition 4, 8
mask 5
masking 4
pipelining work-items on a stream core . . . 4
relationship to work-group 4
relationship with work-groups 4
required number spawned by GPU 4
size 4
size for optimum hardware usage 4
size on AMD GPUs 4
total execution time 5
work-group 2, 3
work-item processing 4
Windows
calling convention 7
compiling
Intel C (C++) compiler 2
OpenCL 2
Visual Studio 2008 Professional Edition 2
debugging
OpenCL kernels 4
running code 6
settings for compiling OpenCL 2
work-group
allocation of LDS size 2
barriers 4
composed of wavefronts 2
concept relating to compute kernels 2
defined 4
dividing global work-size into sub-domains 3
dividing work-items 3
number of wavefronts in 3
performance 2
relationship to wavefront 4
relationship with wavefronts 4
specifying
wavefronts 3
work-item
reaching point (barrier) in the code 2
work-item
branch granularity 4
communicating
through globally shared memory 3
through locally shared memory 3
creation 4
deactivation 9
dividing into work-groups 3

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

element 3
encountering barriers 4
executing
the branch 4
kernel running on compute unit 2
mapping
onto n-dimensional grid (ND-Range) . . . 3
to stream cores 2
non-active 4
processing wavefront 4
reaching point (barrier) in the code 2
scheduling
for execution 3
the range of 2
spawn order 1
synchronization
through barrier operations 3
through fence operations 3
VLIW instruction 4
work-items
divergence in wavefront 4
pipelining on a stream core 4

X

X Window system
using for CL-GL interoperability 8

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.

Copyright © 2013 Advanced Micro Devices, Inc. All rights reserved.